Construction and Recognition of Hyperbolic 3-Manifolds with Geodesic Boundary

Roberto Frigerio   Carlo Petronio

3rd September 2001

Abstract

We extend to the context of hyperbolic 3-manifolds with geodesic boundary Thurston's approach to hyperbolization by means of geometric triangulations. In particular, we introduce moduli for (partially) truncated hyperbolic tetrahedra, and we discuss consistency and completeness equations. Moreover, building on previous work of Ushijima, we extend Weeks' tilt formula algorithm, which computes the Epstein-Penner canonical triangulation, to an algorithm that computes the Kojima triangulation. The theory is particularly interesting in the case of complete finite-volume manifolds with geodesic boundary in which the boundary is non-compact. We include this case using a suitable adjustment of the notion of ideal triangulation, and we show that the case naturally arises within the theory of knots and links.

MSC(2000): 57M50 (primary), 57M25 (secondary).

Contents

Introduction

The aim of this paper is to lay down the theoretical background for a census of orientable hyperbolic 3-manifolds with geodesic boundary. Our starting point is the idea of turning the construction of the hyperbolic structure on a manifold into an algebraic problem. This idea is originally due to Thurston [12] for the case of cusped manifolds, and has been systematically exploited by the software SnapPea [17]. In the setting of cusped manifolds one employs ideal tetrahedra, which are parameterized by complex numbers, and tries to solve the consistency and completeness equations. In the bounded case one has to consider truncated tetrahedra, and moduli get more complicated, but basically the whole scheme extends. The two phenomena of non-compactness and presence of geodesic boundary can actually occur simultaneously, and, following Kojima [8, 9], we introduce the notion of partially truncated tetrahedron to deal with this fact. One interesting point emerges when the boundary of a finite-volume hyperbolic manifold is itself non-compact. Namely, we show that in this case the combinatorial datum to start from to build the structure is not an ideal triangulation of the original manifold, but of a certain quotient of the original manifold.

Working with moduli and equations one can *construct* hyperbolic manifolds with boundary, but, after a list of manifolds has been put together, one has to remove duplicates to get the genuine list, so one is naturally faced with the issue of *recognizing* the manifolds. It turns out that for both cusped and bounded manifolds a certain *canonical* decomposition exists, due to Epstein and Penner [5] in the former case and to Kojima [8, 9] in the latter. One natural strategy to recognize a manifold decomposed into geometric pieces is then to modify the decomposition until the canonical one is reached. This method was described by Weeks in [16] in terms of a so-called "tilt formula", and it was used in [4] in the cusped case. The tilt formula itself was already discussed for the bounded case in [15], and we describe in this paper the whole strategy to turn an arbitrary triangulation of a manifold with boundary into its Kojima decomposition. We warn the reader that, both in the cusped and in the bounded case, the algorithm to transform a decomposition into the canonical one is not proved to converge in general, but, at least in the cusped case, it usually does in practice.

Various differences arise between the cusped and the bounded case, and it is maybe worth mentioning here at least the most subtle one, which requires quite some effort to deal with. Just as the Epstein-Penner decomposition in the cusped case, the Kojima decomposition for bounded manifolds is obtained by projecting to hyperbolic 3-space the faces of a certain polyhedron in Minkowski 4-space. In both cases the polyhedron is the convex hull of certain points that represent, in a suitable sense, liftings of cusps and of boundary components. When there are cusps only, the height of the liftings is intrinsically determined *a priori* (up to global rescaling), and the basic idea to modify a triangulation into the canonical one is to lift the ideal tetrahedra with vertices at the lifted cusps, and make sure the lifted tetrahedra bound a convex set. Essentially the same happens when the boundary is non-empty but there are no cusps at all. In the mixed case, however, only boundary components have a prescribed height to be lifted at, while the height for cusps is a lot harder to determine. This matter is discussed in Sections 4 and 6.

We believe that the issue of understanding and enumerating hyperbolic 3-manifolds with geodesic boundary is a very natural one, and we are planning to exploit the theory developed in this paper in the close future, building an analogue "with boundary" of the cusped census of [4]. Here are three specific reasons for caring about manifolds with boundary:

- These manifolds still satisfy the rigidity theorem, so every geometric invariant, such as the volume or the length spectrum, is actually a topological invariant;
- Thurston's hyperbolization theorem for Haken manifolds [13] implies that all manifolds with boundary satisfying some very natural and fairly general topological properties actually are hyperbolic, so one expects to find that "most" manifolds with boundary are hyperbolic;
- If L is a link in $\mathbb{S}^3$ and Σ is a minimal-genus Seifert surface for L, then the manifold obtained by cutting $\mathbb{S}^3$ along Σ is a natural candidate for a finite-volume hyperbolic structure with boundary. In addition, L has a well-defined "length" with respect to this structure (if any).

Given the length and comparative variety of topics touched in the paper, we have included at the beginning of each section a couple of explanatory paragraphs, where we outline the contents of the section and we list the statements, definitions and notations used later in our work. The reader willing to reach the core of our arguments may at first concentrate on this material only.

1 Triangulations of hyperbolic 3-manifolds with geodesic boundary

In this section we prove some preliminary facts about the topology and geometry at infinity of a finite-volume orientable hyperbolic 3-manifold with geodesic boundary. We also explain what do we mean by a triangulation of such a manifold, showing in particular that this notion must be understood with some care when the boundary of the manifold is non-compact. The essential points of this section are Proposition 1.1, Definitions 1.6, 1.7, 1.8, 1.10, 1.12, 1.13, and 1.14, and Proposition 1.15.

However, Proposition 1.5 and the discussion following it are also quite important as a motivation.

Natural compactification Let N be a complete finite-volume orientable hyperbolic 3-manifold with geodesic boundary. (In the rest of the paper we will summarize all this information saying just that N *is hyperbolic.*) We denote by $D(N)$ the double of N, *i.e.* the manifold obtained by mirroring N in its boundary. Now $D(N)$ is an orientable finite-volume hyperbolic 3-manifold without boundary, so it consists of a compact portion together with several cusps of the form $T \times [0, \infty)$, where T is the torus —see *e.g.* [3]. Within $D(N)$ we have the surface ∂N which cuts $D(N)$ into two isometric copies of N, and to understand the geometry of the ends of N we must investigate how ∂N can intersect a cusp $T \times [0, \infty)$. Using the geometry of $T \times [0, \infty)$ one sees that, up to resizing the cusp, either ∂N is disjoint from $T \times [0, \infty)$ or it is given by $\gamma \times [0, \infty)$, where γ is the union of a finite number of parallel geodesic loops on T. In the first case the cusp $T \times [0, \infty)$ is contained in one of the two isometric copies of N. In the second case, knowing that ∂N is separating in $D(N)$, we see that γ contains at least two loops, and N has an end of the form $A \times [0, \infty)$ where $A \subset T$ is an annulus bounded by these two loops. Since the double of A already is a torus, we also see that γ consists of precisely two loops.

The previous discussion shows that N consists of a compact portion together with some cusps based either on tori or on annuli, which implies the following:

Proposition 1.1. *If N is hyperbolic (*i.e.* N is a complete finite-volume orientable hyperbolic 3-manifold with geodesic boundary) then it has a natural compactification $\overline{N}$ obtained by adding some tori and annuli.*

In particular, ∂N, which we know [8] to be a finite-area orientable hyperbolic surface, can be non-compact. Moreover the ends of ∂N naturally come into pairs $\{\pm 1\} \times S^1 \times [0, \infty) = \partial([-1, 1] \times S^1 \times [0, \infty))$. For later purpose we denote by $\mathcal{A} \subset \overline{N}$ the family of annuli added to compactify N. No specific notation for the tori is needed.

Remark 1.2. If a cusp of N is based on a torus, it is well-known that this torus has a Euclidean structure well-defined up to rescaling. Now, if a cusp is based on an annulus, its double is a Euclidean torus, so the annulus is itself Euclidean with geodesic boundary, up to rescaling. In particular, the annulus is obtained from a Euclidean rectangle by identifying two opposite edges. So, if we normalize the width of the annulus to unity, we can assign the annulus a well-defined *length.*

Topological restrictions We have shown so far that a hyperbolic N is obtained from a compact $\overline{N}$ by removing from $\partial \overline{N}$ some toric components and a family $\mathcal{A}$ of closed embedded annuli. We also know that the components of ∂N are hyperbolic surfaces, whence:

Proposition 1.3. *The components of ∂N have negative Euler characteristic.*

Corollary 1.4. *$\partial\overline{N}$ does not contain spheres, and no annulus of $\mathcal{A}$ can lie on a toric component of $\partial\overline{N}$.*

Proof. There cannot be a sphere because an innermost annulus on a sphere bounds an open disc, having $\chi = 1$. For the same reason on a toric component there cannot be trivial annuli, so there are some parallel annuli, and the complement also consists of annuli, having $\chi = 0$. $\square$

This lemma shows that from the pair $(\overline{N}, \mathcal{A})$ determined by N we can get back N in a non-ambiguous way by removing from $\overline{N}$ both $\mathcal{A}$ and all the toric components of $\partial\overline{N}$. We also have the following additional topological restrictions, stated separately because harder to check directly when an *a priori* non-hyperbolic N is given.

Proposition 1.5. *The compact manifold $\overline{N}$ is irreducible and geometrically atoroidal. Moreover $\overline{N} \setminus \mathcal{A}$ is boundary-incompressible and the only proper essential annuli it contains are parallel in $\overline{N}$ to the annuli in $\mathcal{A}$.*

Proof. Of course $\overline{N}$ is irreducible, because its double $\overline{D(N)}$ is. An embedded incompressible torus must be boundary parallel in $\overline{D(N)}$, whence also in $\overline{N}$.

The toric boundary components of $\overline{N} \setminus \mathcal{A}$ are incompressible because they are in $\overline{D(N)}$. Let Δ be a disc that compresses a loop γ contained in a component Σ of ∂N. Then Δ lifts to the universal cover of N, and the lifting of γ lies on a hyperbolic plane that covers Σ. It readily follows that γ must be trivial in Σ.

An essential annulus cannot join two toric components of $\partial(\overline{N} \setminus \mathcal{A})$, otherwise it would in $\overline{D(N)}$. It also cannot join a toric component with a non-toric one, otherwise its double would join two tori in $\partial\overline{D(N)}$. If an essential annulus joins two non-toric components of $\partial(\overline{N} \setminus \mathcal{A})$ then it lies in N, and its double is an essential torus in $D(N)$. This torus must be boundary-parallel, which easily implies that the annulus is parallel to $\mathcal{A}$ in $\overline{N}$. $\square$

In the previous statement one should notice that irreducibility holds for $\overline{N}$ if and only if it holds for N, and similarly for atoroidality, whereas boundary-incompressibility for $\overline{N} \setminus \mathcal{A}$ does not imply the same property for $\overline{N}$. It is also not possible to deduce from the statement that $\overline{N}$ is anannular. Note however that $\mathcal{A} = \emptyset$ when in N there are no annular cusps.

Links and Seifert surfaces We show in this paragraph that manifolds satisfying (most of) the topological restrictions of Propositions 1.3 and 1.5 naturally arise in the context of the theory of knots and links. Namely, let $L \subset \mathbb{S}^3$ be a link, and let Σ be an orientable Seifert surface for L. Thicken Σ to a product $\Sigma \times [-1,1] \subset \mathbb{S}^3$ so that $\Sigma = \Sigma \times \{0\}$, and define N as $\mathbb{S}^3 \setminus (\Sigma \times (-1,1))$. Note that N compactifies to a manifold $\overline{N}$ by adding the annuli $\mathcal{A} = L \times [-1,1]$ that define the null framing on the components of L. Moreover:

- If $\partial\overline{N}$ contains a sphere then L has a trivial component unlinked from the rest;
- If $\partial\overline{N}$ contains a torus then L contains two parallel components;
- If $\overline{N}$ is not irreducible then L is a split link;
- If $\overline{N}$ is not atoroidal then L is a satellite of a non-trivial knot K and L is homologically trivial in the neighbourhood of K;
- If ∂N is compressible then Σ is the result of a stabilization of another Seifert surface; in particular, Σ cannot have minimal genus.

These remarks provide rather flexible sufficient conditions for $\overline{N}$ to satisfy most of the topological requirements for hyperbolicity. The restriction that essential annuli in N should be parallel to $\mathcal{A}$ in $\overline{N}$ is more involved, and it is not addressed here.

Partially truncated tetrahedra Recall that, when N is finite-volume non-compact hyperbolic and $\partial N = \emptyset$, it is typically possible to decompose N into pieces isometric to geodesic ideal tetrahedra in $\mathbb{H}^3$, and in practice the hyperbolic structure of N is constructed by first taking a topological ideal triangulation and then choosing the geometric shape of the tetrahedra so that their structures match under the gluings giving a complete structure on N. Our wish in the rest this of section is to extend the notion of ideal triangulation to the case of hyperbolic manifolds with geodesic boundary. We begin by describing the pieces into which manifolds will be decomposed, first topologically and then geometrically.

Definition 1.6. We call *partially truncated tetrahedron* a triple $(\Delta, \mathcal{I}, \mathcal{Z})$ where Δ is a tetrahedron, $\mathcal{I}$ is a set of vertices of Δ, and $\mathcal{Z}$ is a set of edges of Δ such that neither of the two endpoints of an edge in $\mathcal{Z}$ belongs to $\mathcal{I}$. The elements of $\mathcal{I}$ and $\mathcal{Z}$ will be called *ideal vertices* and *length*-0 *edges* respectively, for a reason to be explained soon. In the sequel we will always refer to Δ itself as a partially truncated tetrahedron, tacitly implying that certain $\mathcal{I}$ and $\mathcal{Z}$ are also fixed.

Definition 1.7. Given a partially truncated tetrahedron Δ we define its *topological realization* as the space Δ^* obtained by removing from Δ the ideal vertices, the length-0 edges, and small open stars of the non-ideal vertices. We will call *lateral hexagon* and *truncation triangle* the intersection of Δ^* respectively with a face of Δ

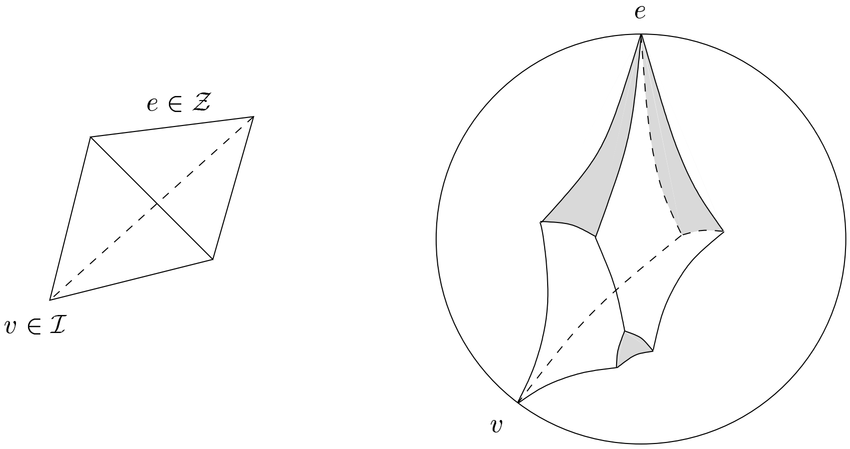

Figure 1: A geometric tetrahedron with one ideal vertex and one length-0 edge.

and with the link in Δ of a non-ideal vertex. The edges of the truncation triangles, which also belong to the lateral hexagons, will be called *boundary edges*. The other edges of the lateral hexagons will be called *internal edges*.

Note that, if Δ has length-0 edges, some vertices of a truncation triangle may be missing. Similarly, if Δ has ideal vertices or length-0 edges, a lateral hexagon of Δ^* may not quite be a hexagon, because some of its (closed) edges may be missing. Note however that two consecutive edges cannot both be missing.

Definition 1.8. Given a partially truncated tetrahedron Δ we call *geometric realization* of Δ an embedding of Δ^* in $\mathbb{H}^3$ such that:

1. The truncation triangles are geodesic triangles, with ideal vertices corresponding to missing vertices;
2. The lateral hexagons are geodesic polygons, with ideal vertices corresponding to missing edges;
3. Truncation triangles and lateral hexagons lie at right angles to each other.

An example of geometric realization is shown in Fig. 1, where truncation triangles are shadowed.

Remark 1.9. If Δ^* is a geometric realization of Δ and v is an ideal vertex of Δ then a neighbourhood of v intersected with Δ^* is automatically isometric in the half-space model $\mathbb{H}^3_{\text{half}} = \mathbb{C} \times (0, \infty)$ of $\mathbb{H}^3$ to $W \times [t_0, \infty)$, where $W \subset \mathbb{C}$ is a triangle and

$t_0 > 0$. Similarly, if e is a length-0 edge then a neighbourhood of e intersected with Δ^* is isometric to $[-1,1]\times[-b,b]\times[t_0,\infty)$. Here the triangles $\{\pm 1\}\times[-b,b]\times[t_0,\infty)$ are contained in the truncation triangles, the triangles $[-1,1]\times\{\pm b\}\times[t_0,\infty)$ are contained in the lateral hexagons, and the closure in Δ of every triangle $\{x\}\times[-b,b]\times[t_0,\infty)$ is obtained by adding only one point of e, so that the segment $[-1,1]\times\{0\}\times\{\infty\}$ can be viewed as a subset of e.

Triangulations In the language introduced above, the classical notion of *ideal triangulation* of a compact 3-manifold with boundary is a realization of the interior of the manifold as a gluing of some Δ^*'s, where the corresponding Δ's have all ideal vertices (and hence no length-0 edge) and the gluing is induced by a simplicial pairing of the faces of the Δ's. We can now easily extend this notion to the situation we are interested in.

Definition 1.10. Let $\overline{N}$ be a compact orientable manifold and let $\mathcal{A}\subset\partial\overline{N}$ be a family of disjoint annuli not lying on the toric components of $\partial\overline{N}$. Let N be obtained from $\overline{N}$ by removing $\mathcal{A}$ and the toric components of $\partial\overline{N}$. We define a *partially truncated triangulation* of N to be a realization of N as a gluing of some Δ^*'s along a pairing of the lateral hexagons induced by a simplicial pairing of the faces of the Δ's.

Remark 1.11.

1. Under the pairing of the faces of the Δ's, ideal vertices are matched to each other. Similarly, length-0 edges are matched to each other.
2. The truncation triangles of the Δ^*'s give a triangulation of ∂N with some genuine and some ideal vertices.
3. The links of the ideal vertices of the Δ's give a triangulation of the toric components of $\partial\overline{N}$.
4. The links of the length-0 edges of the Δ's give a decomposition into rectangles of the annuli in $\mathcal{A}$. On each rectangle $[-1,1]\times[-b,b]$ only the two opposite edges $[-1,1]\times\{\pm b\}$ that lie on lateral hexagons get glued to other rectangles, while the two opposite edges $\{\pm 1\}\times[-b,b]$ that lie on truncation triangles contribute to the boundary of $\mathcal{A}$.

Definition 1.12. Let N as above be endowed with a hyperbolic structure. A partially truncated triangulation of N is called *geometric* if, for each tetrahedron Δ of the triangulation, the pull-back to Δ^* of the Riemannian metric of N defines a geometric realization of Δ. Equivalently, the hyperbolic structure of N should be obtained by gluing geometric realizations of the Δ's along isometries of their lateral hexagons.

In Section 3 we will carefully describe Kojima's result [8] according to which every hyperbolic N as above has a *canonical* decomposition into partially truncated *polyhedra*, rather than tetrahedra. Just as it happens with the Epstein-Penner decomposition [5] of non-compact manifolds with empty boundary, in the vast majority of cases the Kojima decomposition actually consists of tetrahedra, or at least can be subdivided into a geometric partially truncated triangulation.

Manifolds with arcs Our aim is to employ partially truncated triangulations to construct and understand hyperbolic manifolds with boundary, just as ideal triangulations are employed in the cusped case without boundary. One disadvantage of partially truncated triangulations when compared to ideal ones is that the length-0 edges break the symmetry of the tetrahedron, so the situation may appear to be less flexible. It is a useful and remarkable fact that a partially truncated triangulation of a given manifold actually corresponds to a genuine ideal triangulation of another manifold, as we will now explain.

Definition 1.13. Given a manifold N that compactifies to an $\overline{N}$ by adding some tori and a family $\mathcal{A}$ of annuli, we define N' as the quotient of $\overline{N}$ in which every annulus $[-1,1] \times S^1 \in \mathcal{A}$ is collapsed to an arc $[-1,1] \times \{*\}$. Note that N' is also a compact manifold and $[-1,1] \times \{*\}$ is an arc properly embedded in N', as one readily checks by visualizing a neighbourhood of $[-1,1] \times S^1$ as $[-1,1] \times \{z \in \mathbb{C} : 1 \leqslant |z| < 2\}$. We denote by α_N the family of all the arcs $[-1,1] \times \{*\}$ in N'.

Definition 1.14. If M is compact and β is a family of disjoint properly embedded arcs in M, we call *ideal triangulation* of the pair (M, β) an ideal triangulation of M that contains as edges all the arcs in β.

Proposition 1.15. *Partially truncated triangulations of N bijectively correspond to ideal triangulations of (N', α_N).*

Proof. A partially truncated tetrahedron Δ^* is turned into an ideal one by removing the truncation triangles and adding the length-0 edges minus their ends. On the manifold this corresponds to removing the boundary and collapsing each rectangle $[-1,1] \times [-b,b]$ as in Remark 1.11(4) to $[-1,1] \times \{*\}$, see also Remark 1.9. So we get precisely the interior of N' with the arcs in α_N being the length-0 edges.

An ideal triangulation of (N', α_N) is turned into a partially truncated triangulation of N by declaring to be length-0 the edges in α_N and to be ideal the vertices on the tori of $\partial N'$ on which there are no ends of arcs in α_N. □

Having seen how partially truncated triangulations relate to ideal ones, it is natural to ask whether the Matveev-Piergallini calculus [10, 11] for ideal triangulations generalizes to the case of manifolds with arcs. Recall that the fundamental move of this calculus is the two-to-three move (shown below in Fig. 12) which destroys a triangle and the two tetrahedra incident to it, and creates one edge and three tetrahedra incident to this edge. Of course a positive two-to-three move can always be applied to an ideal triangulation of (M, β), while the inverse three-to-two move can be applied as long as the edge it destroys does not lie in β. The next result is due to Amendola (see also [1] and [14]). Since it is not strictly speaking necessary for the present paper, we omit its proof.

Theorem 1.16. *Let two ideal triangulations of the same (M, β) be given. Assume both triangulations contain at least two tetrahedra. Then they are related to each other by a finite combination of two-to-three moves and three-to-two moves that do not destroy the edges of β.*

2 Moduli and equations for partially truncated tetrahedra

In this section we introduce moduli for the geometric realizations of partially truncated tetrahedra, and we describe the equations ensuring that a gluing of geometric tetrahedra gives rise to a consistent and complete hyperbolic structure with geodesic boundary. The idea here is to start with a topological triangulation of a certain manifold with boundary, and try to construct its geometric structure, if any, by choosing the geometric shape of the tetrahedra in the triangulation. We devote the initial part of the section to putting this idea in context and providing motivations, and only then we turn to moduli and equations.

Also in this section we single out the very basic points on which the reader could first concentrate. Moduli are introduced in Theorem 2.2, and consistency equations in Theorem 2.13, with notation coming from formulae (1) to (8) and Fig. 2. Completeness equations are shown to be essentially the same as in the cusped case, and informally discussed after Remark 2.14. In Theorem 2.16 we also show that a solution, if any, is unique.

Hyperbolization with boundary The results of the previous section show that to build a census of hyperbolic 3-manifolds with geodesic boundary we should first list, according to some natural ordering, all the pairs $(\overline{N}, \mathcal{A})$ where $\overline{N}$ is compact and $\mathcal{A} \subset \overline{N}$ is a family of disjoint annuli not lying on the toric components of $\partial\overline{N}$. For each such pair we should then consider the manifold N obtained by removing from $\partial\overline{N}$ all the toric components and the annuli of $\mathcal{A}$, and discard the pair if the conditions of Propositions 1.3 or 1.5 are violated. For each remaining pair we should test the corresponding N for hyperbolicity.

Looking more closely at the strategy just described, one sees that it is very easy to describe an algorithm that lists, *with repetitions*, all the pairs $(\overline{N}, \mathcal{A})$ such that the corresponding N satisfies the condition of Proposition 1.3. The conditions of Proposition 1.5 are harder to check but still manageable, at least theoretically, by means of the technology of normal surfaces. Now, if a pair $(\overline{N}, \mathcal{A})$ survives the topological tests, we see that the double $D(N)$ of the corresponding N has ends of the form $T \times [0, \infty)$ and compactifies to an irreducible, atoroidal, and anannular 3-manifold. Assuming either that N is non-compact or that ∂N is non-empty, we see that $D(N)$ is Haken, so Thurston's hyperbolization theorem [13] shows that $D(N)$ is finite-volume hyperbolic, and the involution of $D(N)$ that fixes ∂N and interchanges N with its mirror copy can be realized by an isometry [6], so N also is hyperbolic. However, this theoretical proof of existence of the hyperbolic structure is not satisfactory under at least two respects. First, it does not allow to compute the geometric invariants of N, such as the volume. Second, it leaves unsettled the issue of removing duplicates from the list of manifolds.

The alternative strategy based on triangulations which we will now describe overcomes both the drawbacks of the topological approach just pointed out. It should be noted, however, that *a priori* there could exist hyperbolic manifolds that cannot be triangulated geometrically. These manifolds would be missed by our search.

Enumeration strategy To employ triangulations, we switch from the $(\overline{N}, \mathcal{A})$ to the (N', α_N) compactification of the candidate hyperbolic N. So, our first step is to list all pairs $(\mathcal{T}, \alpha)$ where $\mathcal{T}$ is an ideal triangulation of some compact orientable 3-manifold N' with boundary, and α is a set of edges of $\mathcal{T}$, also viewed as a set of properly embedded arcs in N'. Of course there are infinitely many such $(\mathcal{T}, \alpha)$'s, so in practice one always deals with some finite "initial" segment of the list. A pair $(\mathcal{T}, \alpha)$ is immediately discarded if $\partial N'$ contains spheres on which there are two or fewer ends of the arcs in α. If $(\mathcal{T}', \alpha)$ is not discarded, we define $\overline{N}$ as N' minus an open tubular neighbourhood for each arc in α, and $\mathcal{A}$ as the family of annuli that bound the removed tubes. Now we can define N as $\overline{N}$ minus $\mathcal{A}$ and the boundary tori, and N automatically satisfies the condition of Proposition 1.3.

A pair $(\mathcal{T}, \alpha)$ gives rise to a partially truncated triangulation of the corresponding N by declaring to be length-0 the edges in α, and to be ideal the vertices corresponding to the toric components of $\partial \overline{N}$ on which there is no end of any arc in α. We will prove in the rest of this section that there exists an algorithm to answer the question whether *can the tetrahedra of $\mathcal{T}$ be geometrically realized in $\mathbb{H}^3$ so to define a complete hyperbolic structure on N*. If the answer is affirmative then we add N to our census, if not we pass to the next $(\mathcal{T}, \alpha)$.

Remark 2.1. To make the search more effective, a slight modification of the method just described can be employed. Namely, when the shape of the elements of a triangulation $\mathcal{T}$ of a certain (N', α) cannot be chosen to give a complete structure on the corresponding N, it is often convenient, before giving up, to try with other triangulations of the same (N', α). It typically happens, at least in the non-compact empty-boundary case dealt with by SnapPea [17], that eventually a triangulation is found that either is geometric or suggests which of the topological restrictions of Proposition 1.5 is violated.

The outcome of the strategy just outlined is a list of hyperbolic manifolds, each with a certain geometric triangulation. However this list contains repetitions, that we can remove if we can recognize manifolds. Concentrating on those which have non-empty boundary, we note that each of them has a unique well-defined Kojima decomposition. This decomposition can now be viewed as the *name* of the manifold, because two such decompositions can be checked to be equal or not by comparing the geometric shape of the polyhedra and the combinatorics of the gluings. The recognition issue is then reduced to the issue of constructing the Kojima decomposition starting from an arbitrary geometric triangulation. This is the theme we will concentrate on starting from the next section.

Moduli We will now show that the dihedral angles at the non-0-length edges can be used as moduli for geometric tetrahedra. As explained below in Remark 2.8, for a tetrahedron without ideal vertices, the lengths of the internal edges could also be employed, but of course not in general.

Theorem 2.2. *Let Δ be a partially truncated tetrahedron and let $\Delta^{(1)}$ be the set of edges of Δ. The geometric realizations of Δ are parameterized up to isometry by the functions $\theta : \Delta^{(1)} \to [0, \pi)$ such that:*

- $\theta(e) = 0$ *if and only if e is length-0;*
- *For each vertex v of Δ, if e_1, e_2, e_3 are the edges that emanate from v, then $\theta(e_1) + \theta(e_2) + \theta(e_3)$ is equal to π for ideal v and less than π for non-ideal v.*

The map θ corresponding to a geometric realization Δ^* associates to each non-0-length edge e the dihedral angle $\theta(e)$ of Δ^* along e.

Proof. Our argument follows Fujii's [7]. Let $\theta : \Delta^{(1)} \to [0, \pi)$ be as in the statement. We fix notation as in Fig. 2 and set $\theta_i = \theta(e_i)$. Our task is to show that there exists and is unique up to isometry a geometric realization Δ^* of Δ with dihedral angles θ_i along the e_i's. The idea is to construct the four planes in $\mathbb{H}^3$ on which the faces should lie, and to prove that their configuration is unique and determines Δ^* up to

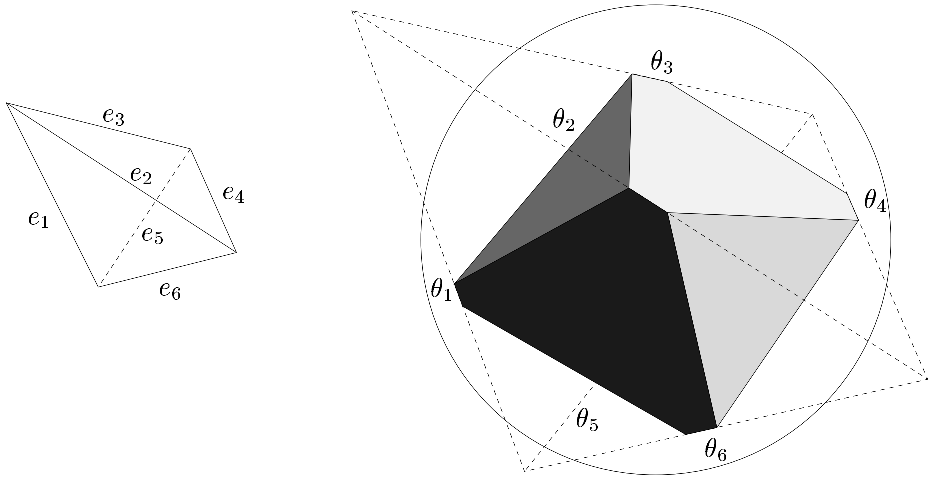

Figure 2: Notation for edges and dihedral angles of a truncated tetrahedron. The geometric figure on the right uses the projective model of $\mathbb{H}^3$, as explained below in Section 3.

isometry. The plane containing the face with edges e_i, e_j, e_k will be determined by its circle $C_{ijk} \subset \partial\mathbb{H}^3$ of points at infinity. We use the half-space model $\mathbb{H}^3_{\text{half}}$ with $\partial\mathbb{H}^3_{\text{half}} = \mathbb{E}^2 \cup \{\infty\}$ and we identify a line $\ell \subset \mathbb{E}^2$ to the circle $\ell \cup \{\infty\}$.

We first assume that Δ has neither ideal vertices nor length-0 edges, and we show that the configuration of the C_{ijk}'s exists and is determined by the θ_i's. Later we will prove that the configuration determines a unique Δ^*. We choose C_{126} and C_{135} to be lines through $0 \in \mathbb{E}^2$ at angle θ_1 with each other. Conditions $\theta_1 + \theta_2 + \theta_3 < \pi$ and $\theta_1 + \theta_5 + \theta_6 < \pi$ imply quite easily that there exist circles C_{234} and C_{456} as in Fig. 3. Next, we modify C_{456} by a dilation and use conditions $\theta_3 + \theta_4 + \theta_5 < \pi$ and $\theta_2 + \theta_4 + \theta_6 < \pi$ to show that it can be placed as in Fig. 4. This shows that the configuration of the C_{ijk}'s exists. Its uniqueness follows from uniqueness up to rotation and dilation of the configuration of lines and circles of Fig. 4. Now we show that the C_{ijk}'s determine Δ^* uniquely. The truncation plane at the vertex v_{123} (where e_1, e_2, e_3 have their common end) must be orthogonal to $C_{126}, C_{135}, C_{234}$. From Fig. 3 we see that such a plane exists and is unique, and similarly for the other three truncation planes. Moreover the truncation triangles are pairwise disjoint, and the conclusion follows.

The same argument applies with minor changes to the case where Δ has ideal vertices or length-0 edges. See for instance Fig. 6 below for the configuration to use in case $\theta_1 + \theta_2 + \theta_3 = \pi$ and $\theta_6 = 0$. $\square$

Corollary 2.3. *The space of geometric realizations of Δ up to isometry is a bounded*

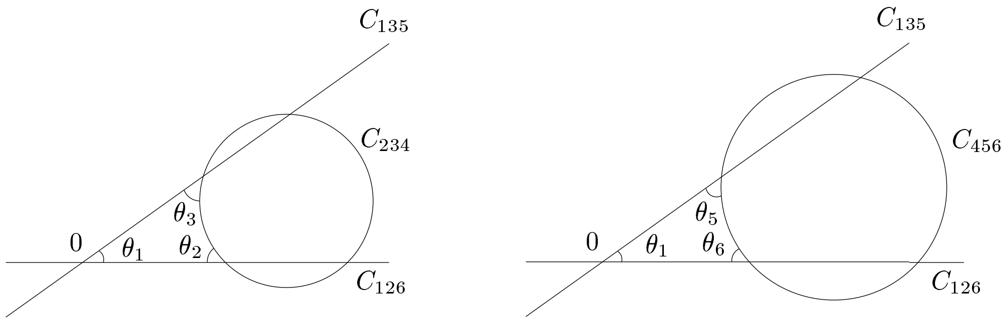

Figure 3: Lines and circles at prescribed angles with each other.

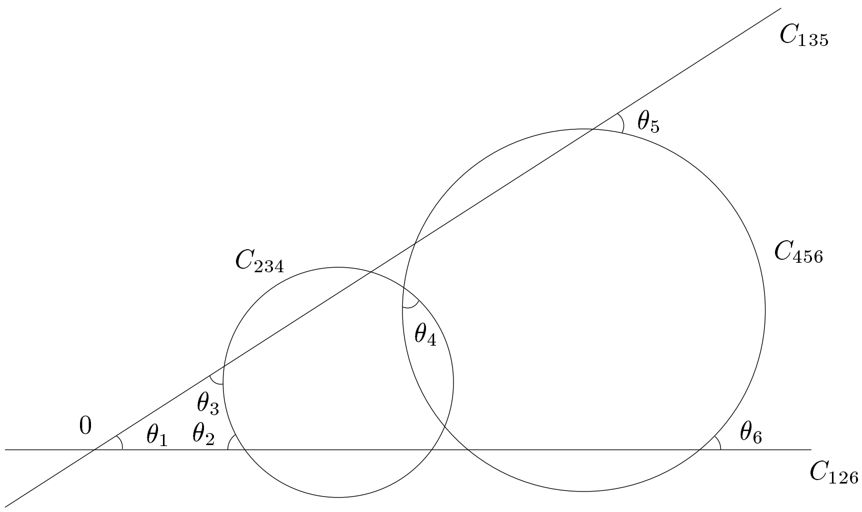

Figure 4: Trace at infinity of a geometric realization.

open subset of a real Euclidean space of dimension $6-n$ where n is the total number of ideal vertices and length-0 edges of Δ.

From angles to lengths Having introduced moduli for geometric tetrahedra, our next task is to determine, given a triangulated orientable manifold, which choices of moduli for the tetrahedra give a global hyperbolic structure on the manifold. There are two obvious necessary conditions (which are often but not always sufficient, *e.g.* they are not when all the vertices are ideal). Namely, we should have a total dihedral angle of 2π around each non-0-length edge of the manifold, and we should be able to glue the lateral hexagons by isometries. The first condition is directly expressed in terms of moduli. To express the second condition recall that the shape of a hyperbolic right-angled hexagon is determined by the lengths of a triple of pairwise disjoint edges. This may seem to suggest that, to ensure consistency, one only has to compute, in terms of the dihedral angles, either the lengths of the internal edges or the lengths of the boundary edges. This is however false when ideal vertices and/or length-0 edges are involved, so we will need to compute both.

Let us consider a partially truncated tetrahedron Δ with edges $e_1,\ldots,e_6$ as above in Fig. 2. In the rest of this paragraph we fix a geometric realization θ of Δ determined by dihedral angles $\theta_i=\theta(e_i)$ for $i=1,\ldots,6$, and we denote by L^θ the length with respect to this realization. The boundary edges of the lateral hexagons of Δ correspond to the pairs of distinct non-opposite edges $\{e_i,e_j\}$, and will be denoted by e_{ij}. Now e_{ij} disappears towards infinity, so it has length 0, when the common vertex of e_i and e_j is ideal, it is an infinite half-line when one of e_i or e_j is 0-length, and it is an infinite line when both e_i and e_j are 0-length. The next result, that is readily deduced from [2, The Cosine Rule II, pag. 148] allows to compute the length of e_{ij} when this length is finite. We refer to e_{12} with notation as in Fig. 2.

Proposition 2.4. *If both e_1 and e_2 have non-0 length then*

$$\cosh L^\theta(e_{12})=\frac{\cos\theta_1\cdot\cos\theta_2+\cos\theta_3}{\sin\theta_1\cdot\sin\theta_2}. \tag{1}$$

Note that this result is correct (and obvious) also when the common end of e_1 and e_2 is ideal. Turning to the length of an internal edge, we note that the edge is an infinite half-line or an infinite line when one or both its ends are ideal. Otherwise the length is computed using [2, The Cosine Rule II, pag. 148 and Theorem 7.19.2, pag. 161]. To state the result of the computation we need to introduce certain functions that will be used again below. With notation as in Fig. 2, and defining

v_{ijk} as the vertex from which the edges e_i, e_j, e_k emanate, we set:

$$\begin{aligned} c^\theta(e_1) &= \cos\theta_1 \cdot (\cos\theta_3 \cdot \cos\theta_6 + \cos\theta_2 \cdot \cos\theta_5) \\ &\quad + \cos\theta_2 \cdot \cos\theta_6 + \cos\theta_3 \cdot \cos\theta_5 + \cos\theta_4 \cdot \sin^2\theta_1; \end{aligned} \tag{2}$$

$$d^\theta(v_{123}) = 2\cos\theta_1 \cdot \cos\theta_2 \cdot \cos\theta_3 + \cos^2\theta_1 + \cos^2\theta_2 + \cos^2\theta_3 - 1. \tag{3}$$

Lemma 2.5. *$d^\theta(v_{123}) = 0$ if and only if the vertex v_{123} is ideal.*

Proposition 2.6. *If v_{123} and v_{156} are both non-ideal then*

$$\cosh L^\theta(e_1) = c^\theta(e_1) \Big/ \sqrt{d^\theta(v_{123}) \cdot d^\theta(v_{156})}. \tag{4}$$

The next fact will be proved in Section 6 using results from Sections 3 and 4.

Proposition 2.7. *If Δ has no ideal vertices then a geometric realization of Δ^* is determined up to isometry by the lengths of its internal edges.*

Remark 2.8. The previous proposition implies that, when there are no ideal vertices, one could employ the lengths of the internal edges as moduli. Besides the loss of generality, this choice is however inadvisable because of the following drawbacks:

- In terms of lengths, the restriction that the three dihedral angles at each vertex should sum up to less than π gets replaced by somewhat more complicated relations. Namely, one should express boundary lengths in terms of internal lengths, and then for each vertex impose the triangular inequalities for the three boundary edges at that vertex;

- Dihedral angles are needed in any case to ensure consistency, and to express angles in terms of lengths one should invert formula (4), which does not appear to be completely straight-forward.

Exceptional hexagons Recall that we are looking for the conditions to ensure that a gluing of lateral hexagons can be realized by isometries. By default gluings match ideal vertices to each other and length-0 edges to each other, because these notions are part of the initial topological information about a triangulation.

When a pairing glues two *compact* hexagons, *i.e.* when there are no ideal vertices or length-0 edges involved, to make sure that the gluing is an isometric one we may equivalently require the lengths of the internal edges or those of the boundary edges to match under the gluing. If we actually require *all* lengths to match, we are guaranteed that the gluing is isometric also for non-compact hexagons, except in the very special case where a boundary edge disappears into an ideal vertex, and the opposite internal edge is length-0. In this case, which is illustrated in Fig. 5,

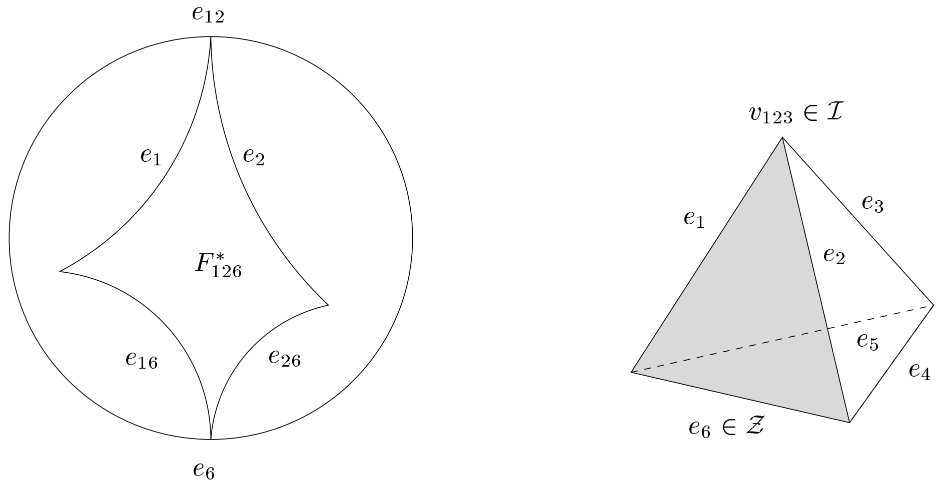

Figure 5: A hexagon with a boundary edge and the opposite internal edge being length-0.

there is no length at all to match, because two edges have length 0, and the other four are infinite half-lines. But two hexagons as in Fig. 5 need not be isometric to each other, as the next discussion shows.

To parameterize the special hexagons we need now to be slightly more careful about orientation than we have been so far. Namely, we choose on the tetrahedra an orientation compatible with a global orientation of the manifold. As a result also the lateral hexagons and the truncation triangles have a fixed orientation, and the gluing maps reverse the orientation of the hexagons.

So, let us consider an exceptional hexagon F^*_{126} as in Fig. 5, with the boundary edge e_{12} lying at an ideal vertex, and the opposite internal edge e_6 being length-0. Recall that the hexagon is oriented and embedded in $\mathbb{H}^3$ by θ. We consider now the horospheres O_1 and O_2 centred at e_{12} and passing through the non-ideal ends of e_1 and e_2 respectively. We define $\sigma^\theta(F_{126})$ to be $\pm\text{dist}(O_1, O_2)$, the sign being positive if e_2, e_{12}, e_1 are arranged positively on ∂F^*_{126} and O_1 is contained in the horoball bounded by O_2, or if e_2, e_{12}, e_1 are arranged negatively on ∂F^*_{126} and O_2 is contained in the horoball bounded by O_1, and negative otherwise. This definition easily implies the following:

Proposition 2.9. *Let F and F' be paired exceptional lateral hexagons. Their pairing can be realized by an isometry if and only if $\sigma^\theta(F) + \sigma^\theta(F') = 0$.*

Proposition 2.10. *Let F_{126} be the exceptional hexagon of Fig. 5, oriented so that*

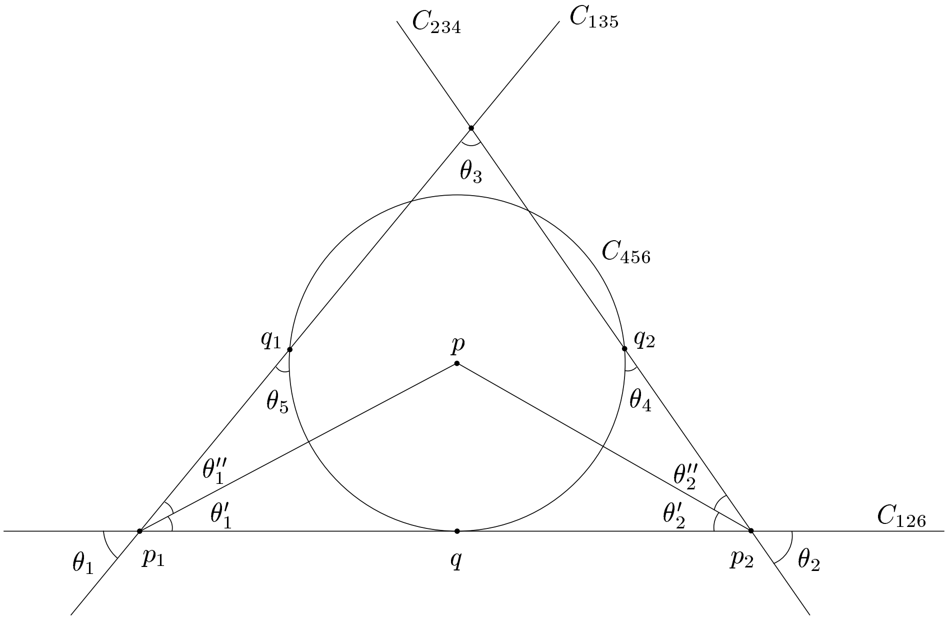

Figure 6: Configuration at infinity when there is an exceptional hexagon.

e_2, e_{12}, e_1 are positively arranged on ∂F^*_{126}. Then

$$\sigma^\theta(F_{126}) = \log \frac{\sin\theta_2}{\cos\theta_2 + \cos\theta 4} - \log \frac{\sin\theta_1}{\cos\theta_1 + \cos\theta_5}. \tag{5}$$

Proof. We realize Δ^* in $\mathbb{H}^3_{\text{half}}$ so that $v_{123} = \infty$ and denote by C_{ijk} the circle at infinity of the plane that contains F^*_{ijk}. Let $C_{126} \cap C_{135} = \{p_1, \infty\}$ and $C_{126} \cap C_{234} = \{p_2, \infty\}$. Condition $\theta_6 = 0$ means that C_{456} is tangent to C_{126}. The configuration is then as in Fig. 6, where we also introduce more notation needed for the proof. The truncation plane relative to v_{156} is now the Euclidean half-sphere of radius $L(qp_1)$ centred at p_1, so the finite end of e_1 has coordinates $(p_1, L(qp_1))$. Similarly the finite end of e_2 is $(p_2, L(qp_2))$, so $\sigma^\theta(F_{126}) = \log L(qp_1) - \log L(qp_2) = \log L(qp_1)/L(qp_2)$. Of course $L(qp_1)/L(qp_2) = \tan\theta_2'/\tan\theta_1'$.

Now $\angle(pq_1p_1) = \theta_5 + \pi/2$, whence $\sin\angle(pq_1p_1) = \cos\theta_5$, and the sine theorem yields $L(pq_1)/L(pp_1) = \sin\theta_1''/\cos\theta_5$. But $L(pq_1) = L(pq)$, so $L(pq_1)/L(pp_1) = \sin\theta_1'$. Equaling the two expressions of $L(pq_1)/L(pp_1)$ we get $\sin\theta_1'' = \cos\theta_5 \cdot \sin\theta_1'$. On the other hand $\theta_1'' = \theta_1 - \theta_1'$, so $\sin\theta_1'' = \sin\theta_1 \cdot \cos\theta_1' - \cos\theta_1 \cdot \sin\theta_1'$. Equaling the two expressions of $\sin\theta_1''$ and dividing by $\cos\theta_1'$ we get $\tan\theta_1' = \sin\theta_1/(\cos\theta_1 + \cos\theta_5)$. Similarly $\tan\theta_2' = \sin\theta_2/(\cos\theta_2 + \cos\theta_4)$ and the conclusion follows. □

Matching around edges We have discussed so far the conditions under which the hyperbolic structure of the geometric tetrahedra matches across lateral hexagons. As already mentioned, there is another obvious condition we must impose if we want the structure to extend also along the internal edges of the triangulation. Namely, let us define for an edge e

$$\alpha^\theta(e) = \sum\{\theta(e') : \ e' \text{ is glued to } e\}. \tag{6}$$

Then $\alpha^\theta(e)$ should be 2π for all e. This condition is actually sufficient when there are no ideal vertices, but not in general. The point is that when the geometric tetrahedra are arranged one after each other around a non-0-length edge e, the first face of the first tetrahedron and the second face of the last tetrahedron may overlap without coinciding. Namely, the isometry which pairs these two faces may be a translation along e instead of being the identity. Of course the isometry has to be the identity unless both ends of e are ideal.

We recall now that a horospherical cross-section near an ideal vertex v of a geometric tetrahedron is a Euclidean triangle well-defined up to similarity. The tetrahedron being oriented, this triangle is also oriented, so, once a vertex of the triangle is fixed, its similarity structure is determined by a complex parameter in the upper half-plane π_+. Choosing a vertex of the triangle amounts to choosing an edge e ending at v, so we have a well-defined modulus $z^\theta(e,v) \in \pi_+$ whenever v is ideal and e ends at v. We also define

$$Z^\theta(e,v) = \prod\{z^\theta(e',v') : \ (e',v') \text{ is glued to } (e,v)\}. \tag{7}$$

The next result is proved just as in the purely ideal case (see [3]):

Proposition 2.11. *Assume the structure defined by θ matches across lateral hexagons, and let e be an internal edge with both ideal ends. If v is any one of these ends, the structure matches across e if and only if $\alpha^\theta(e) = 2\pi$ and $Z^\theta(e,v) = 1$.*

Proposition 2.12. *For an oriented tetrahedron as in Fig. 2, assume v_{123} is ideal and e_1, e_2, e_3 are positively arranged around v_{123}. Then:*

$$z^\theta(e_1, v_{123}) = \frac{\sin\theta_2}{\sin\theta_3} \cdot e^{i\theta_1}. \tag{8}$$

Consistency equations We have preferred above to introduce bit after bit our conditions for the geometric tetrahedra to define a global structure on the manifold, but now we collect the relevant information in one precise statement.

Theorem 2.13. *Consider an orientable manifold N obtained from a compact $\overline{N}$ by removing all the tori and some annuli contained in $\partial\overline{N}$. Suppose that all the components of ∂N have negative Euler characteristic. Fix a partially truncated triangulation of N, with the tetrahedra oriented so that the gluings reverse the induced orientations. Let θ be a geometric realization of the tetrahedra in the triangulation. Then θ defines on N a hyperbolic structure with geodesic boundary if and only if the following conditions hold:*

1. $L^\theta(e) = L^\theta(e')$ *for all pairs* (e, e') *of matching internal edges;*
2. $L^\theta(e) = L^\theta(e')$ *for all pairs* (e, e') *of matching boundary edges;*
3. $\sigma^\theta(F) + \sigma^\theta(F') = 0$ *for all pairs* (F, F') *of matching exceptional hexagons;*
4. $\alpha^\theta(e) = 2\pi$ *for all edges* e*;*
5. $Z^\theta(e, v) = 1$ *for all edges* e *with both ideal ends and for both ends* v *of* e.

Remark 2.14. By Propositions 2.4, 2.6, 2.10, and 2.12, all the above conditions can be expressed as analytic equations in terms of the dihedral angles θ. Moreover:

- If N has no toric end, *i.e.* if $\partial\overline{N}$ contains no tori, then condition (2) is a consequence of (1), and (3) and (5) are empty. So, to ensure hyperbolicity of N, one may impose (1) and (4) only;
- If N has no annular end, *i.e.* if no annuli are removed from $\partial\overline{N}$, then condition (1) is a consequence of (2), and (3) is empty;
- When conditions (1), (2), and (3) are in force, condition (5) may be equivalently imposed at either end of the edge e.

Completeness equations The discussion of completeness is very easy. If N is hyperbolic, N is complete if and only if its double $D(N)$ is. So we should ensure the toric ends of $D(N)$ to be complete, *i.e.* their bases to have an induced a Euclidean structure, rather than just a similarity structure. If an end of $D(N)$ comes from a toric cusp of N, completeness is imposed as usual by requiring the holonomy of the similarity structure on the torus to consist of translations. If an end of $D(N)$ is the double of an annular end of N then it is actually always complete. To see this, we use Remarks 1.9 and 1.11, which imply that the annulus at the basis of the end of N is tiled by a (cyclic) row of Euclidean rectangles. So the annulus is Euclidean and its boundary circles have the same length, whence the double is a Euclidean torus.

We remind the reader that completeness of an end $T \times [0, \infty)$ can be turned into a pair of equations in terms of the moduli of the triangular horospheric cross-sections, and hence in terms of the dihedral angles by means of Proposition 2.12. To do this, one first chooses as a basis of $H_1(T;\mathbb{Z})$ a pair of loops which are simplicial with respect to the triangulation of T induced by the triangulation of N. Then one notes that the dilation component of the holonomy of a simplicial loop γ is the product of all moduli γ leaves to its left, multiplied by -1 if γ has an odd number of vertices —see *e.g.* [3].

Uniqueness A crucial fact for computational purposes is that a solution of the hyperbolicity equations, if any, is unique. Before showing this we spell out the rigidity theorem already mentioned above.

Theorem 2.15. *On any given manifold there exists, up to isometry, at most one finite-volume and complete hyperbolic structure with geodesic boundary.*

Proof. If N is finite-volume hyperbolic with boundary then $D(N)$ is finite-volume hyperbolic without boundary, so its structure is unique, and we only need to show that the embedding of ∂N in $D(N)$ cannot be homotoped away from itself. Note that ∂N, being geodesic, is automatically incompressible. Uniqueness of embedding for ∂N now easily follows by considering non-trivial simple loops on ∂N and recalling that every non-trivial free-homotopy class of a loop in $D(N)$ has a unique geodesic representative. □

Theorem 2.16. *Given a triangulated manifold N as in Theorem 2.13, there exists at most one choice of the geometric realization θ that turns N into a complete manifold with geodesic boundary.*

Proof. Assume θ_0 and θ_1 yield complete structures on N. By the rigidity theorem these structures are actually the same, so we can view θ_0 and θ_1 as defining isotopic geometric triangulations of one hyperbolic N. Of course two geometric triangulations are identical if they have the same edges, so we are left to show that if $(e_t)_{t\in[0,1]}$ is an isotopy of properly embedded segments, half-lines, or lines, and e_0, e_1 are geodesic, then $e_0 = e_1$. Consider first the case of a segment. If we double N we get for all t a closed loop $D(e_t)$ in $D(N)$. The free-homotopy class of $D(e_t)$ is of course independent of t, and it must be non-trivial, otherwise e_0 would lift in $\mathbb{H}^3$ to a geodesic segment with both ends on a component of the lifting of ∂N, so e_0 would actually be contained in ∂N. So all the $D(e_t)$ lift to infinite open lines in $\mathbb{H}^3$, and these lines have two well-defined ends. These ends are fixed points of hyperbolic isometries from a discrete group, and they evolve continuously along the isotopy. This implies that the ends are actually independent of t, whence $e_0 = e_1$.

The same argument applies when one or both the ends of the edges e_t tend to infinity along a cusp, except that hyperbolic fixed points get replaced by parabolic fixed points. □

Remark 2.17. As a by-product of the previous argument we deduce that if a triangulation contains a boundary-parallel non-0-length edge, then the triangulation is never geometric.

3 The Kojima decomposition

In this section we describe the canonical decomposition due to Kojima [8] of a hyperbolic 3-manifold with non-empty geodesic boundary, recalling several details because we will be using them throughout the rest of the paper. We omit all proofs addressing the reader to [8]. All the notation introduced in this section is employed extensively later on, so there is basically nothing the reader could skip here.

Minkowsky space Kojima's construction takes place in 4-dimensional Minkowsky space, so we start by fixing some notation about it. We denote by $\mathbb{M}^{3+1}$ the space $\mathbb{R}^4$ with coordinates x_0, x_1, x_2, x_3 endowed with the Lorentzian inner product $\langle x, y\rangle = -x_0y_0 + x_1y_1 + x_2y_2 + x_3y_3$. We set

$$\begin{aligned} \mathcal{H}^3_- &= \{x \in \mathbb{M}^{3+1} : \langle x, x\rangle = -1,\ x_0 > 0\}, \\ \mathcal{H}^3_+ &= \{x \in \mathbb{M}^{3+1} : \langle x, x\rangle = 1\}, \\ L^3_+ &= \{x \in \mathbb{M}^{3+1} : \langle x, x\rangle = 0,\ x_0 > 0\}. \end{aligned}$$

We recall that $\mathcal{H}^3_-$ is the upper sheet of the two-sheeted hyperboloid, and that $\langle\,\cdot\,,\,\cdot\,\rangle$ restricts to a Riemannian metric on $\mathcal{H}^3_-$. With this metric, $\mathcal{H}^3_-$ is the so-called *hyperboloid model* $\mathbb{H}^3_{\rm hyp}$ of hyperbolic space. The one-sheeted hyperboloid $\mathcal{H}^3_+$ turns out to have a bijective correspondence with the set of hyperbolic half-spaces in $\mathbb{H}^3_{\rm hyp}$. Given $w \in \mathcal{H}^3_+$, the corresponding half-space, called the *dual* of w, is given by

$$\{v \in \mathcal{H}^3_- : \langle v, w\rangle \leqslant 0\}.$$

Similarly, the cone L^3_+ of future-oriented light-like vectors of $\mathbb{M}^{3+1}$ corresponds to the set of horospheres in $\mathbb{H}^3_{\rm hyp}$. The horosphere dual to $u \in L^3_+$ is given by

$$\{v \in \mathcal{H}^3_- : \langle v, u\rangle = -1\}.$$

Projective model and truncated polyhedra Let $\pi : \mathbb{M}^{3+1} \setminus \{0\} \to \mathbb{P}(\mathbb{M}^{3+1})$ be the canonical projection. We set $\Pi^3 = \{x \in \mathbb{M}^{3+1} : x_0 = 1\}$ and note that Π^3 can be viewed as a subset of $\mathbb{P}(\mathbb{M}^{3+1})$ via π. Moreover Π^3 is isometric to Euclidean 3-space $\mathbb{E}^3$, and π restricts to a bijection between $\mathcal{H}^3_-$ and the unit ball of Π^3. Giving this ball the metric that turns π into an isometry with $\mathbb{H}^3_{\rm hyp}$, we get the projective model $\mathbb{H}^3_{\mathrm{proj}}$ of hyperbolic space. This model is particularly suitable for describing partially truncated tetrahedra, as we will now show.

Take in Π^3 a tetrahedron Δ with distinct vertices outside $\mathbb{H}^3_{\mathrm{proj}}$ or on its boundary, and call *ultra-ideal* the vertices not lying on $\partial\mathbb{H}^3_{\mathrm{proj}}$ (those on $\partial\mathbb{H}^3_{\mathrm{proj}}$ are called *ideal* as usual). Suppose that the interior of every edge of Δ meets $\mathbb{H}^3_{\mathrm{proj}}$ or $\partial\mathbb{H}^3_{\mathrm{proj}}$ at least in one point. Each ultra-ideal vertex v of Δ determines a hyperbolic plane H_v in $\mathbb{H}^3_{\mathrm{proj}}$ obtained as the intersection of $\mathbb{H}^3_{\mathrm{proj}}$ with the Lorentzian orthogonal to v. Let us now define Δ^* as the intersection of Δ with $\mathbb{H}^3_{\mathrm{proj}}$ truncated by the planes H_v (see Fig. 2 above). It is easy to see that Δ^* is a geometric realization of the partially truncated tetrahedron Δ in which the ideal vertices are those on $\partial\mathbb{H}^3_{\mathrm{proj}}$, and the length-0 edges are those tangent to $\partial\mathbb{H}^3_{\mathrm{proj}}$. This implies in particular that our notation Δ^* is consistent.

Before proceeding it is worth noting that the hyperbolic plane H_v described above can also be constructed by elementary Euclidean geometry on Π^3. Namely, if we take the cone in $\Pi^3 \cong \mathbb{E}^3$ with vertex v and tangent to $\partial\mathbb{H}^3_{\mathrm{proj}} \cong \mathbb{S}^2$, then H_v is the Euclidean disc bounded by the circle where the cone intersects $\partial\mathbb{H}^3_{\mathrm{proj}}$.

The above definition of Δ^* of course makes sense also for convex polyhedra Δ more complicated than tetrahedra, provided Δ only has ideal and ultra-ideal vertices, and the interior of every edge of Δ meets the closure of $\mathbb{H}^3_{\mathrm{proj}}$. Partially truncated polyhedra of this sort are the blocks of the Kojima decomposition described in the rest of this section.

Convex hull of lifted boundary components Let N be hyperbolic with non-empty geodesic boundary. Identifying the universal cover of the double $D(N)$ of N with $\mathbb{H}^3$, we can realize the universal cover of N itself as a closed convex region $\widetilde{N}$ of $\mathbb{H}^3$ bounded by a locally finite countable family $\mathcal{S}$ of pairwise disjoint planes. The group of deck transformations of the covering $\widetilde{N} \to N$, denoted henceforth by Γ, is the stabilizer of $\widetilde{N}$ in the group of deck transformations of $\mathbb{H}^3 \to D(N)$. In the rest of the section we will always use the $\mathbb{H}^3_{\mathrm{proj}}$ model of $\mathbb{H}^3$.

Noting that $\widetilde{N}$ lies on a definite side of each plane $S \in \mathcal{S}$, we consider now the vector of $\mathcal{H}^3_+$ dual to the half-plane that contains $\widetilde{N}$ and is bounded by S. We denote by $\mathcal{B}$ the family of all these duals, and we define $\mathcal{C} \subset \mathbb{M}^{3+1}$ as the closure of $\mathrm{Conv}(\mathcal{B})$, where $\mathrm{Conv}(X)$ denotes from now on the convex hull of a set $X \subset \mathbb{M}^{3+1}$. Kojima has shown that $\mathcal{B}$ is a discrete subset of $\mathbb{M}^{3+1}$, that $0 \notin \mathcal{C}$ and that $\pi(\mathcal{C}) \supset \mathbb{H}^3_{\mathrm{proj}}$. In particular $\mathcal{C}$ has non-empty interior, and it is Γ-invariant by construction. The idea is now to construct a Γ-equivariant tessellation of $\widetilde{N}$ that projects to a decomposition of N by intersecting $\widetilde{N}$ with the projections to $\mathbb{H}^3_{\mathrm{proj}}$ of the 3-faces of $\partial\mathcal{C}$. However, it turns out that not all faces should be projected, and that some faces have non-trivial stabilizer in Γ, so they must be subdivided. To explain the matter in detail we begin with the following:

Definition 3.1. Let X be a subset of $\mathbb{M}^{3+1}$ such that $0 \notin X$. A point $x \in X$ is called *almost-visible* with respect to X if the segment $[0, x]$ meets X in x only. The point is called *visible* if it is almost-visible and $\pi(x) \in \mathbb{H}^3_{\mathrm{proj}}$.

Of course only the faces of $\partial\mathcal{C}$ containing visible points should contribute to the tessellation of $\widetilde{N}$. These faces are also called *visible*. It turns out that there are two quite different sorts of visible faces, called respectively *elliptic* and *parabolic* depending on whether the restriction of $\langle\cdot,\cdot\rangle$ to the hyperplane on which the face lies is positive-definite or positive-semi-definite.

Cut locus and elliptic faces The first type of visible faces of $\partial\mathcal{C}$ correspond to the vertices of the cut-locus of $\widetilde{N}$ relative to $\partial\widetilde{N}$, which we now define for an arbitrary manifold M.

Definition 3.2. Let M be hyperbolic with non-empty geodesic boundary. We define the *cut-locus* $\mathrm{Cut}(M, \partial M)$ of M relative to ∂M as the set of points of M that admit at least two different shortest paths to ∂M. A point is called a *vertex* of the cut-locus if it admits four different shortest paths to ∂M whose initial tangent vectors span the tangent space to M at the point as an affine space.

The next result is implicit in Kojima's work [8]. A proof is readily deduced from Proposition 4.10 shown below (using discreteness, which is easy to establish).

Proposition 3.3. *$\mathrm{Cut}(N, \partial N)$ has finitely many vertices. A point of $\widetilde{N}$ is a vertex of $\mathrm{Cut}(\widetilde{N}, \partial\widetilde{N})$ if and only it projects in N to a vertex of $\mathrm{Cut}(N, \partial N)$.*

For every vertex v of $\mathrm{Cut}(\widetilde{N}, \partial\widetilde{N})$ we define now $\mathcal{B}(v) \subset \mathcal{B}$ as the set of dual vectors to the hyperplanes in $\mathcal{S}$ having shortest distance from v. Since in $\mathbb{H}^3$ there is a unique shortest path joining a given point to a given plane, each $\mathcal{B}(v)$ contains at least four vectors. The next result describes the visible elliptic faces of $\partial\mathcal{C}$.

Proposition 3.4. *For every vertex v of $\mathrm{Cut}(\widetilde{N}, \partial\widetilde{N})$ there exists a unique elliptic hyperplane $E(v) \subset \mathbb{M}^{3+1}$ such that $E(v) \cap \mathcal{B} = \mathcal{B}(v)$. Moreover $E(v)$ is a support hyperplane for $\mathcal{C}$, and:*

1. *The points in $E(v) \cap \mathcal{C}$ are almost-visible with respect to $\mathcal{C}$;*
2. *$E(v) \cap \mathcal{C}$ is a 3-dimensional compact polyhedron whose set of vertices is $\mathcal{B}(v)$;*
3. *The stabilizer of $E(v) \cap \mathcal{C}$ in Γ is trivial.*

Parabolic faces and subdivision Besides those corresponding to vertices of the cut-locus, $\partial\mathcal{C}$ has visible faces coming from toric cusps of N. Let us denote by $\widetilde{N}_\infty \subset \partial\mathbb{H}^3_{\mathrm{proj}}$ the set of points at infinity of $\widetilde{N}$. A point $q \in \widetilde{N}_\infty$ is said to *generate a toric cusp* if it is fixed under a $\mathbb{Z}\oplus\mathbb{Z}$ subgroup Γ_q of parabolic elements of Γ, *i.e.* if there is a horoball centred at q that projects to a toric cusp of N.

For $\tilde{q}$ in L^3_+ and $t \in \mathbb{R}$ we consider now the affine parabolic hyperplane

$$F(\tilde{q}, t) = \{v \in \mathbb{M}^{3+1} : \langle v, \tilde{q}\rangle = -t\}.$$

Proposition 3.5. *Let $q \in \widetilde{N}_\infty$ generate a toric cusp of N, and take $\tilde{q} \in L^3_+$ such that $\pi(\tilde{q}) = q$. Then there exists a unique $t(\tilde{q}) \in \mathbb{R}$ such that $F(\tilde{q}, t(\tilde{q}))$ is a support hyperplane for $\mathcal{C}$. Moreover $F(\tilde{q}, t(\tilde{q}))$ depends only on q, not on $\tilde{q}$, and, setting $F(q) = F(\tilde{q}, t(\tilde{q}))$, we have that:*

1. *$t(\tilde{q}) > 0$ and $\mathcal{C} \subseteq \bigcup_{t\geqslant t(\tilde{q})} F(\tilde{q}, t)$;*
2. *The points in $F(q) \cap \mathcal{C}$ are almost-visible with respect to $\mathcal{C}$;*
3. *$F(q) \cap \mathcal{B}$ is infinite and $F(q) \cap \mathcal{C}$ is a non-compact 3-dimensional polyhedron whose set of vertices is $F(q) \cap \mathcal{B}$;*
4. *The 2-dimensional faces of $F(q) \cap \mathcal{C}$ are compact;*
5. *The stabilizer in Γ of $F(q) \cap \mathcal{C}$ coincides with the $\mathbb{Z}\oplus\mathbb{Z}$ stabilizer Γ_q of q.*

Point (5) of this proposition shows that $F(q)\cap\mathcal{C}$ must be subdivided before projecting to $\mathbb{H}^3_{\mathrm{proj}}$ and intersecting with $\widetilde{N}$. Continuing with the same notation, for every 2-dimensional face W of $F(q) \cap \mathcal{C}$ we define now $W^{(\tilde{q})}$ as the cone based on W with vertex in $\tilde{q}$. Note that $W^{(\tilde{q})}$ depends on $\tilde{q}$, and it meets $\mathcal{C}$ in W only. However one easily sees that $W^{(q)} = \pi(W^{(\tilde{q})})$ actually does not depend on $\tilde{q}$. Moreover the family of all $W^{(q)}$'s, as W varies in the 2-faces of $F(q) \cap \mathcal{C}$, gives a Γ_q-equivariant tessellation of $\{\tilde{q}\} \cup \pi(F(q) \cap \mathcal{C})$ in which every polyhedron has trivial stabilizer.

Canonical decomposition We begin with the following fact:

Proposition 3.6. *The faces of $\mathcal{C}$ described in Propositions 3.4 and 3.5 contain all the visible points of $\mathcal{C}$.*

In addition to this, one can easily show that if Δ is a visible face of $\mathcal{C}$ then the (partial) truncation $\pi(\Delta)^*$, defined earlier in this section, of $\pi(\Delta) \subset \Pi^3$ is obtained by intersection with $\widetilde{N}$.

We are now ready to summarize the construction. Let us denote by $\mathcal{V}$ the set of vertices of $\mathrm{Cut}(\widetilde{N}, \partial\widetilde{N})$, and by Q the family of all points of $\widetilde{N}_\infty$ that generate toric cusps of N. For each $q \in Q$ we fix an arbitrary $\tilde{q} \in L^3_+$ such that $\pi(\tilde{q}) = q$, and we denote by $\widetilde{Q}$ the family of all such $\tilde{q}$'s. We define $\mathcal{K}_{\mathbb{M}}(\widetilde{Q})$ to be the family of all visible elliptic faces $E(v) \cap \mathcal{C}$ and all faces $W^{(\tilde{q})}$ obtained from the visible parabolic faces. Here v varies in $\mathcal{V}$, $\tilde{q}$ varies in $\widetilde{Q}$, and W varies in the 2-dimensional faces of $F(\pi(\tilde{q})) \cap \mathcal{C}$. We also denote by $\mathcal{K}_{\mathbb{P}}$ the family of polyhedra obtained by projecting to Π^3 the elements of $\mathcal{K}_{\mathbb{M}}(\widetilde{Q})$, and by $\mathcal{K}^*_{\widetilde{N}}$ the family obtained by intersecting with $\widetilde{N}$ (or, equivalently, truncating) the elements of $\mathcal{K}_{\mathbb{P}}$. We know that indeed $\mathcal{K}_{\mathbb{P}}$ and $\mathcal{K}^*_{\widetilde{N}}$ are independent of $\widetilde{Q}$, and we have:

Theorem 3.7. *$\mathcal{K}^*_{\widetilde{N}}$ is a Γ-equivariant decomposition of $\widetilde{N}$ into partially truncated polyhedra. Every element of $\mathcal{K}^*_{\widetilde{N}}$ has trivial stabilizer in Γ, so $\mathcal{K}^*_{\widetilde{N}}$ projects to a canonical and finite decomposition $\mathcal{K}^*_N$ of N into partially truncated polyhedra.*

Remark 3.8. By construction, each polyhedron in $\mathcal{K}_{\mathbb{P}}$ has at most one ideal vertex. All other vertices are ultra-ideal.

4 Choice of heights

The general strategy to decide whether a given decomposition of a hyperbolic N is the canonical one is as in [16], *i.e.* to lift the decomposition first to $\mathbb{H}^3_{\mathrm{proj}}$ and then to $\mathbb{M}^{3+1}$, and to make sure that the resulting polyhedra bound a convex set. In the setting of the construction described in the previous section, the lifting of a polyhedron to $\mathbb{M}^{3+1}$ can be performed directly when N does not have toric cusps, because one only has to select which points of $\mathcal{B}$ are the vertices of the lifting, and $\mathcal{B}$ depends on N only. When there are toric cusps, however, we have an arbitrariness for the liftings of the ideal vertices, corresponding to the arbitrariness of the choice of $\widetilde{Q}$. In addition, for some choices of $\widetilde{Q}$ it can happen that $\mathcal{K}_{\mathbb{M}}(\widetilde{Q})$ does *not* bound a convex set. The case of toric cusps must therefore be discussed with some care.

The reader willing to catch just the main points of the section and proceed to the sequel could devote attention only to Definition 4.1, formula (9), Definition 4.4, and Propositions 4.7, 4.9, 4.10, and 4.11.

Horospherical cross-sections We first show how to reduce the choice of $\widetilde{Q}$ to a choice that is more intrinsic, but still arbitrary at this stage. We fix N hyperbolic with $\partial N \neq \emptyset$ for the rest of the section.

Definition 4.1. Let N be hyperbolic. A family $\mathcal{O}$ of disjoint tori in N that lift to horospheres in $\widetilde{N}$ and bound disjoint cusps of N will be called a *horospherical cross-section* of N. Given such an $\mathcal{O}$, we define $\widetilde{Q}(\mathcal{O})$ as the lifting of Q where a point q is lifted to the only $\tilde{q} \in L^3_+$ such that $\pi(\tilde{q}) = q$ and the dual to $\tilde{q}$ projects in N to a component of $\mathcal{O}$. We also define $\mathcal{K}_{\mathbb{M}}(\mathcal{O})$ as $\mathcal{K}_{\mathbb{M}}(\widetilde{Q}(\mathcal{O}))$.

Remark 4.2. The cross-section $\mathcal{O}$ is determined by a sufficiently small positive number assigned to each toric cusp, namely the volume of the region between the cross-section at the cusp and infinity. Insisting all volumes to be equal to each other, one may also determine $\mathcal{O}$ by a single number. This number is naturally interpreted as the (inverse of the) *height* at which the cross-section should be taken —the smaller the volume, the higher the cross-section.

Augmented convex hull Before proceeding, we need to define a set of which $\mathcal{K}_{\mathbb{M}}(\mathcal{O})$ gives the visible boundary. Namely, we set:

$$\mathcal{C}'(\mathcal{O}) = \mathcal{C} \cup \bigcup \Big\{ \mathrm{Conv}\Big(\{\tilde{q}\} \cup (F(q) \cap \mathcal{C}) \Big) : \ q \in Q, \ \tilde{q} = \widetilde{Q}(\mathcal{O}) \cap \pi^{-1}(q) \Big\}. \tag{9}$$

The fact that indeed the visible boundary of $\mathcal{C}'(\mathcal{O})$ is $\mathcal{K}_{\mathbb{M}}(\mathcal{O})$ easily follows from point (1) of Proposition 3.5.

Lemma 4.3. *Suppose x is an almost-visible point of $\mathcal{C}'(\mathcal{O})$ and $\pi(x) \in \widetilde{N}$. Then $\alpha \cdot x \in \mathcal{C}'(\mathcal{O})$ for every $\alpha \geqslant 1$.*

Proof. The statement for $\mathcal{C}'(\mathcal{O})$ easily follows from the corresponding statement for $\mathcal{C}$. In addition, since $\mathcal{C}$ is closed, it is sufficient to prove the statement when $\pi(x)$ lies in the interior of $\widetilde{N}$. Convexity of $\widetilde{N}$ then implies that $\pi(x)$ lies in the interior of the convex hull of a finite number of points of $\widetilde{N}_\infty$. Moreover we know from [8] that $\widetilde{N}_\infty = \overline{\pi(\mathcal{B})} \cap \partial\mathbb{H}^3_{\mathrm{proj}}$, and it is easy to deduce that $\pi(x)$ lies in the convex hull of infinitely many finite and pairwise disjoint subsets of $\pi(\mathcal{B})$. But $\mathcal{B}$ is discrete in $\mathbb{M}^{3+1}$, so there is a sequence α_n of reals diverging to $+\infty$ such that $\alpha_n \cdot x \in \mathcal{C}$ for all n. Convexity of $\mathcal{C}$ now implies the desired conclusion. □

Definition 4.4. Let P_1, P_2 be finite 3-dimensional convex polyhedra in $\mathbb{M}^{3+1} \setminus \{0\}$ which project injectively to $\mathbb{P}(\mathbb{M}^{3+1})$. Assume P_1 and P_2 share a 2-face F, and for $i = 1, 2$ let H_i be the half-hyperplane in $\mathbb{M}^{3+1}$ such that $H_i \supset P_i$ and $\partial H_i \supset F$. We say P_1 and P_2 form a *convex angle* at F if the connected component not containing 0 of $\mathbb{M}^{3+1} \setminus (H_1 \cup H_2)$ is convex. The angle is called *strictly convex* if it is convex and $H_1 \cup H_2$ is not a hyperplane.

Lemma 4.5. *The following facts are equivalent:*

1. *$\mathcal{C}'(\mathcal{O}) \cap \pi^{-1}(\widetilde{N})$ is convex;*
2. *The 3-faces of $\mathcal{C}'(\mathcal{O})$ meeting $\pi^{-1}(\widetilde{N})$ form convex angles with each other.*

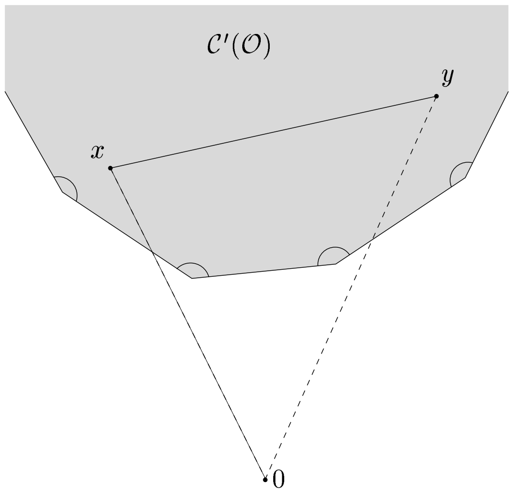

Figure 7: Convexity of $\mathcal{C}'(\mathcal{O}) \cap \pi^{-1}(\widetilde{N})$.

Proof. Implication (1) $\Rightarrow$ (2) is clear. To show the opposite implication, we take $x, y \in \mathcal{C}'(\mathcal{O})$ such that $\pi(x), \pi(y) \in \widetilde{N}$. Using the fact that angles are convex, we want to show that $[x, y] \subset \mathcal{C}'(\mathcal{O})$. Since $\mathcal{C}'(\mathcal{O})$ is closed we can also slightly perturb x and y, so we assume that $\pi(x) \neq \pi(y)$ and $[\pi(x), \pi(y)]$ does not meet the 1-skeleton of the decomposition $\mathcal{K}^*_{\widetilde{N}}$ of $\widetilde{N}$. So $[\pi(x), \pi(y)]$ meets finitely many 2-faces, and the picture on $\pi^{-1}([x, y])$ is as in Fig. 7 by the assumption on convexity of angles. The conclusion now follows from Lemma 4.3. $\square$

The previous result easily implies that if $\mathcal{C}'(\mathcal{O}) \cap \pi^{-1}(\widetilde{N})$ is convex for some choice of "heights" $\mathcal{O}$, then it is also convex for any "higher" choice. We will now show that such a choice is possible. To this end we will need the following easy result (points (1) and (2) of which are actually only used in Section 5 below):

Lemma 4.6. *1. Let H_1 and H_2 be disjoint planes in $\mathbb{H}^3_{\mathrm{proj}}$. Choose disjoint half-spaces bounded by these planes and let $w_1, w_2 \in \mathcal{H}^3_+$ be the duals to these half-spaces. Then* $\cosh d(H_1, H_2) = -\langle w_1, w_2 \rangle$;

2. *Let W_1 and W_2 be half-spaces in $\mathbb{H}^3_{\mathrm{proj}}$ such that $\partial W_1 \cap \partial W_2 \neq \emptyset$ and let θ be the dihedral angle they determine; let $w_1, w_2 \in \mathcal{H}^3_+$ be the duals to W_1 and W_2. Then $\cos\theta = -\langle w_1, w_2 \rangle$;*

3. *Let O and H be respectively a horosphere and a plane in $\mathbb{H}^3_{\mathrm{proj}}$. Assume $O \cap H = \emptyset$, let $u \in L^3_+$ be the dual to O and $w \in \mathcal{H}^3_+$ be the dual to the subspace bounded by H that contains O. Then* $\exp d(O, H) = -\langle u, w \rangle$;

4. *Let O_1 and O_2 be disjoint horospheres with different centres in $\mathbb{H}^3_{\mathrm{proj}}$, and let $u_1, u_2 \in L^3_+$ be their dual vectors. Then $\exp d(O_1, O_2) = -\frac{1}{2}\langle u_1, u_2\rangle$.*

Proposition 4.7. *Let N be hyperbolic with universal covering $\mathbb{H}^3_{\mathrm{proj}} \supset \widetilde{N} \to N$. Let $\mathcal{O}$ be a horospherical cross-section for N such that:*

1. *For any distinct components O_1 and O_2 of the lifting of $\mathcal{O}$ we have*
$$\exp d(O_1, \partial\widetilde{N}) + \exp d(O_2, \partial\widetilde{N}) < 2 \cdot \exp d(O_1, O_2);$$
2. *The toric cusps in N determined by $\mathcal{O}$ do not contain vertices of $\mathrm{Cut}(N, \partial N)$, and for any such vertex u we have $\sinh d(u, \partial N) < \exp d(u, \mathcal{O})$.*

Then the visible 3-faces of $\mathcal{C}'(\mathcal{O})$ form strictly convex angles with each other. In particular, $\mathcal{C}'(\mathcal{O}) \cap \pi^{-1}(\widetilde{N})$ is convex.

Proof. Recall that we have two types of faces, the elliptic ones $E(v) \cap \mathcal{C}$ where v is a vertex of $\mathrm{Cut}(\widetilde{N}, \partial\widetilde{N})$, and the $W^{(\tilde{q})}$'s where $\tilde{q} \in \widetilde{Q}(\mathcal{O})$ and W is a 2-face of the parabolic face $F(\pi(\tilde{q})) \cap \mathcal{C}$. By construction elliptic faces form strictly convex angles with each other, and the same happens for faces $W_1^{(\tilde{q})}$ and $W_2^{(\tilde{q})}$ relative to the same $\tilde{q}$. So we have two cases to discuss.

We begin with the case of two faces of the form $W_1^{(\tilde{q}_1)}$ and $W_2^{(\tilde{q}_2)}$, which of course can only have a common 2-face when $W_1 = W_2 =: W$. Knowing that the horospheres O_1 and O_2 dual to $\tilde{q}_1$ and $\tilde{q}_2$ are disjoint, it is not hard to see that there exists a real number $a > 1$ and an isometry of $\mathbb{M}^{3+1}$ that carries $\tilde{q}_1$ and $\tilde{q}_2$ to the points $(a, a, 0, 0)$ and $(a, -a, 0, 0)$ respectively. Using Lemma 4.6 (4) it is also easy to see that a is intrinsically interpreted as the square root of $\exp d(O_1, O_2)$.

Of course convexity and all the relevant quantities are preserved under isometry, so we can just assume $\tilde{q}_1 = (a, a, 0, 0)$ and $\tilde{q}_2 = (a, -a, 0, 0)$. Recall now that for $i = 1, 2$ we have in $\mathbb{M}^{3+1}$ the hyperplanes $F(\pi(\tilde{q}_i)) = F(\tilde{q}_i, t(\tilde{q}_i))$, and the face W at which we must prove convexity lies in the intersection of these hyperplanes. Moreover Lemma 4.6 (3) implies that $t(\tilde{q}_i)$ has the intrinsic meaning of $\exp d(O_i, \partial\widetilde{N})$. Now the 2-plane on which the face W lies is the following one:
$$\Big\{x \in \mathbb{M}^{3+1} \,:\, x_0 = \big(t(\tilde{q}_1) + t(\tilde{q}_2)\big) / 2a,\ x_1 = \big(-t(\tilde{q}_1) + t(\tilde{q}_2)\big) / 2a\Big\}.$$

Therefore the angle at W is strictly convex if and only if $(1/2a) \cdot (t(\tilde{q}_1) + t(\tilde{q}_2)) < a$, and this inequality holds by the first assumption of the statement and the intrinsic interpretation of the $t(\tilde{q}_i)$'s and a.

Turning to the angle between a face $W^{(\tilde{q})}$ and a face $E(v) \cap \mathcal{C}$, we denote by O the horosphere dual to $\tilde{q}$, and note that the horoball bounded by O cannot contain

v by assumption. Using this fact it is not hard to show that up to isometry in $\mathbb{M}^{3+1}$ we can assume that $v = (1, 0, 0, 0)$ and $\tilde{q} = (b, b, 0, 0)$ for some $b > 1$. Again using the above lemma, one computes b to be $\exp d(v, O)$. Moreover $E(v)$ is the hyperplane of equation $x_0 = a$, with $a = \sinh d(v, \partial\widetilde{N})$. Now the angle at W is strictly convex if and only if $a < b$, and the conclusion follows. $\square$

Corollary 4.8. *$\mathcal{O}$ can be chosen so that both the above conditions hold.*

Proof. Finiteness of the number of vertices of $\mathrm{Cut}(N, \partial N)$ readily implies that the second condition of the proposition is satisfied for some suitably high $\mathcal{O}$. Now let $\alpha = \max\{d(O, \partial N) :\ O \in \mathcal{O}\}$, and redefine $\mathcal{O}$ by pushing up each horospherical cross-section at distance α from its previous position. A straight-forward computation then implies that also the first condition of the proposition is fulfilled. $\square$

Recognizing the canonical decomposition Assume now that we have some decomposition $\mathcal{D}^*$ of N into partially truncated polyhedra, and let us ask ourselves whether $\mathcal{D}^*$ coincides with Kojima's canonical $\mathcal{K}^*_N$. We fix as above the universal cover $\widetilde{N} \to N$ with $\widetilde{N} \subset \mathbb{H}^3_{\mathrm{proj}}$ and a horospherical cross-section $\mathcal{O}$ of N. Now we can first lift $\mathcal{D}^*$ to a decomposition $\mathcal{D}^*_{\widetilde{N}}$ of $\widetilde{N}$, then we can consider the associated non-truncated family of polyhedra $\mathcal{D}_{\mathbb{P}}$ in Π^3, and finally we can lift $\mathcal{D}_{\mathbb{P}}$ to a $\mathcal{D}_{\mathbb{M}}(\mathcal{O})$ in $\mathbb{M}^{3+1}$ by lifting each ultra-ideal vertex to its representative in $\mathcal{B}$ and each ideal vertex $q \in Q$ to $\widetilde{Q}(\mathcal{O}) \cap \pi^{-1}(q)$.

Proposition 4.9. *Let $\mathcal{D}^*$ be a geometric decomposition of a hyperbolic N and let $\mathcal{O}$ be a horospherical cross-section for N as in Proposition 4.7. Let $\mathcal{D}_{\mathbb{M}}(\mathcal{O})$ be the lifting of $\mathcal{D}^*$ to $\mathbb{M}^{3+1}$ just described. Then:*

- *$\mathcal{D}^*$ is the canonical decomposition $\mathcal{K}^*_N$ if and only if the 3-faces of $\mathcal{D}_{\mathbb{M}}(\mathcal{O})$ form strictly convex angles with each other;*
- *$\mathcal{D}^*$ is a subdivision of the canonical decomposition $\mathcal{K}^*_N$ if and only if the 3-faces of $\mathcal{D}_{\mathbb{M}}(\mathcal{O})$ form convex angles with each other.*

Proof. The "only if" assertions are obvious. To prove the "if" assertions, assume the angles are convex. Define $Y \subset \mathbb{M}^{3+1}$ as the union of the polyhedra of $\mathcal{D}_{\mathbb{M}}(\mathcal{O})$ and set $X = \{\alpha \cdot y :\ \alpha \geqslant 1,\ y \in Y\}$. Using convexity of angles, the same argument given for Lemma 4.5 shows that X is convex. Now $X \supset Y \supset \mathcal{B} \cup \widetilde{Q}(\mathcal{O})$, so $X \supset \mathcal{C}'(\mathcal{O})$ by convexity. Moreover, the polyhedra of $\mathcal{D}_{\mathbb{M}}(\mathcal{O})$ have vertices in $\mathcal{B} \cup \widetilde{Q}(\mathcal{O})$, so $Y \subset \mathcal{C}'(\mathcal{O})$. Since Y is the visible boundary of X, we can conclude that it is also the visible boundary of $\mathcal{C}'(\mathcal{O})$, so $\mathcal{K}_{\mathbb{M}}(\mathcal{O})$ is the natural decomposition of Y into faces. Our given decomposition $\mathcal{D}_{\mathbb{M}}(\mathcal{O})$ of Y then coincides with $\mathcal{K}_{\mathbb{M}}(\mathcal{O})$ when all its angles are strictly convex. If there are flat angles, however, $\mathcal{D}_{\mathbb{M}}(\mathcal{O})$ is only a subdivision of $\mathcal{K}_{\mathbb{M}}(\mathcal{O})$. $\square$

We have now some remarks about how to apply Proposition 4.9 in practice. To begin, note that to ensure (strict) convexity at all the infinitely many 2-faces of $\mathcal{D}_{\mathbb{M}}(\mathcal{O})$, it is actually sufficient to check it for one lifting of each 2-face of $\mathcal{D}$, and there are finitely many of such faces. However, two serious issues remain. First, we need an effective method to check convexity, and we will provide one in Section 5. Second, we need a way to determine $\mathcal{O}$ using the geometry of $\mathcal{D}^*$ only. A partial step in this direction is discussed in the rest of this section, and the conclusion is given (when $\mathcal{D}^*$ is actually a triangulation) in Section 6.

Intrinsic computation of height Now we show how to find a horospherical cross-section as in Proposition 4.7 in terms of the geometry of N.

Proposition 4.10. *Let N be hyperbolic and realize the universal cover $\widetilde{N} \to N$ with $\widetilde{N} \subset \mathbb{H}^3_{\text{half}}$ so that ∞ generates a toric cusp of N. Let Γ_∞ be the $\mathbb{Z} \oplus \mathbb{Z}$ stabilizer of ∞ in the group of deck transformations of $\widetilde{N} \to N$. Let $r_1 > r_2$ be the first and second largest radii of the components of $\partial\widetilde{N}$, viewed as Euclidean half-spheres in $\mathbb{C} \times (0, \infty)$. Let d be the diameter of a fundamental domain for the action of Γ_∞ on $\mathbb{C} \times \{0\}$, and define $k(r_1, r_2, d) = \sqrt{3} \cdot (r_1^2 + d^2/4)/(r_1 - r_2)$. Consider the horosphere $\widetilde{O} = \mathbb{C} \times \{k(r_1, r_2, d)\} \subset \mathbb{H}^3_{\text{half}}$ and let O be the projection of $\widetilde{O}$ in N. Then:*

1. *O is an embedded toric cross-section of the cusp, and it is disjoint from any other arbitrarily chosen embedded toric cross-section at any of the other cusps;*
2. *The cusp bounded by O does not contain vertices of* $\text{Cut}(N, \partial N)$*;*
3. *If u is a vertex of* $\text{Cut}(N, \partial N)$ *we have* $\sinh d(u, \partial N) < \exp d(u, O)$.

Proof. The constants r_1, r_2, d are fixed, so we set $k = k(r_1, r_2, d)$. Since $k > r_1$, we see that $\widetilde{O} \cap \partial\widetilde{N} = \emptyset$. So both assertions of point (1) follow from the following claim: *If $\widetilde{O}'$ is a horosphere centred at some point of $\partial\mathbb{H}^3_{\text{half}} \setminus \{\infty\}$ and $\widetilde{O}' \cap \partial\widetilde{N} = \emptyset$ then $\widetilde{O}' \cap \widetilde{O} = \emptyset$.* To prove this claim, let $\widetilde{O}'$ be a Euclidean sphere of radius x tangent to $\mathbb{C} \times \{0\}$ at a point z. We must show that $2x < k$. Now within distance $d/2$ from z there exists the centre w of a component of $\partial\widetilde{N}$ of Euclidean radius r_1. Knowing that $\widetilde{O}'$ and this component are disjoint, we deduce that $x < (d^2/4 - r_1^2)/(2 \cdot r_1)$ whence the conclusion at once.

To prove points (2) and (3) we set $h = \sqrt{(r_1^2 + d^2/4)/(1 - r_2/r_1)}$ and claim the following: $\mathbb{C} \times [h, \infty)$ *does not contain any vertex of* $\text{Cut}(\widetilde{N}, \partial\widetilde{N})$. Using Proposition 3.3 and the easy fact that $k \geqslant h$, our claim readily implies point (2), and it will be used below for point (3).

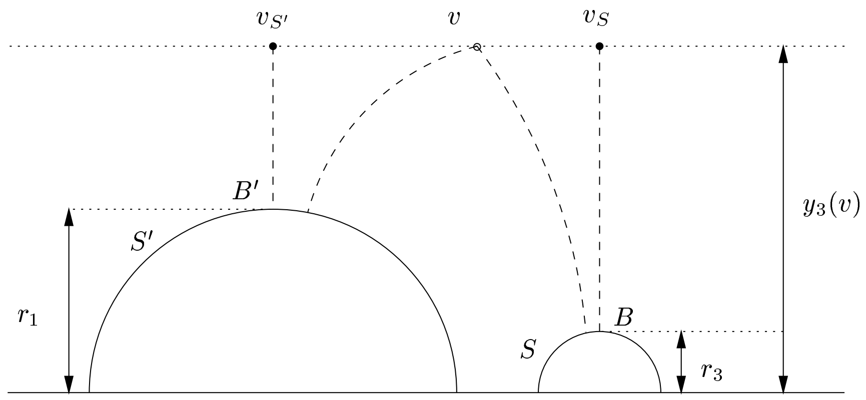

Figure 8: Notation for the proof of Proposition 4.10 (2,3).

To prove the claim, let v a vertex of $\mathrm{Cut}(\widetilde{N}, \partial\widetilde{N})$. We first show that the components of $\partial\widetilde{N}$ nearest to v cannot all have the same Euclidean radius. If this were the case, turning to the setting of Kojima's construction in $\mathbb{M}^{3+1}$ and using Lemma 4.6 (3), we would deduce that $\mathcal{B}(v)$ is contained in a parabolic affine hyperplane of $\mathbb{M}^{3+1}$ (Lorentz orthogonal to the dual of $\widetilde{O}$ in L^3_+), against Proposition 3.4.

Since the components of $\partial\widetilde{N}$ nearest to v do not all have the same radius, one of them, say S, has some radius $r_3 \leqslant r_2$. By definition of d there exists another component S' of $\partial\widetilde{N}$ with radius r_1 such that, with notation as in Fig. 8, the Euclidean distance between v and $v_{S'}$ is at most $d/2$. Recall now that if y is the $(0,\infty)$-coordinate on $\mathbb{H}^3_{\mathrm{half}} = \mathbb{C} \times (0,\infty)$ and $\|\cdot\|$ is the Euclidean norm, the hyperbolic distance between $p, q \in \mathbb{H}^3_{\mathrm{half}}$ satisfies

$$\cosh d(p,q) = 1 + \|p-q\|^2 \,/\, 2y(p)y(q). \tag{10}$$

Now $d(v_S, B) = d(v_S, S) \leqslant d(v, S)$, and we know that $d(v,S) \leqslant d(v,S') \leqslant d(v,B')$, so $d(v_S, B) \leqslant d(v, B')$. Using (10) we easily deduce that

$$y(v)^2 \leqslant \frac{r_1 \cdot r_3 + d^2/4 \cdot r_3/r_1 - r_3^2}{1 - r_3/r_1} \leqslant \frac{r_1^2 + d^2/4}{1 - r_3/r_1} \leqslant \frac{r_1^2 + d^2/4}{1 - r_2/r_1} = h^2.$$

Our claim, and hence point (2), are proved. For point (3) it is sufficient to show that for any vertex v of $\mathrm{Cut}(\widetilde{N}, \partial\widetilde{N})$ we have $\sinh d(v, \partial\widetilde{N}) < \exp d(v, \widetilde{O})$. Choose again S' as in Fig. 8. The claim just proved and the easy fact that $h > r_1$ show that $\|v - B'\|^2 < h^2 + d^2/4$. But $h > d/2$, so $\|v - B'\|^2 < 2h^2$. This inequality and (10) easily imply that

$$\sinh^2 d(v, \partial\widetilde{N}) \leqslant \sinh^2 d(v, B') < \left(h^4 + 2 \cdot h^2 \cdot y(v) \cdot r_1\right) / \left(y(v)^2 \cdot r_1^2\right).$$

As noted above we have $r_1 < h$. Moreover $y(v) < h$ by the claim shown, so

$$\sinh^2 d(v, \partial\widetilde{N}) < (3 \cdot h^4) \,/\, (y(v)^2 \cdot r_1^2) = k^2 \,/\, y(v)^2 = \exp^2 d(v, \widetilde{\mathcal{O}}),$$

and the proof is complete. □

Proposition 4.11. *Let a hyperbolic N have n toric cusps, and for $i = 1, \ldots, n$ choose a realization $\widetilde{N}_i$ of $\widetilde{N}$ in $\mathbb{H}^3_{\mathrm{half}}$ so that ∞ generates the i-th cusp. Let $r_1^{(i)}, r_2^{(i)}, d^{(i)}$ be the constants relative to $\widetilde{N}_i$ as in Proposition 4.10. Set*

$$\lambda = \max\left\{k\left(r_1^{(i)}, r_2^{(i)}, d^{(i)}\right)/r_1^{(i)} \,:\, i = 1, \ldots, n\right\}.$$

Define O_i as the projection of $\mathbb{C} \times \{\lambda \cdot k(r_1^{(i)}, r_2^{(i)}, d^{(i)})\}$ from $\widetilde{N}_i$ to N. Then $\{O_i\}_{i=1}^n$ is a horospherical cross-section as in Proposition 4.7.

Proof. If O_i is first defined as the projection of $\mathbb{C} \times \{k(r_1^{(i)}, r_2^{(i)}, d^{(i)})\}$, then $d(O_i, \partial N)$ equals $\log(k(r_1^{(i)}, r_2^{(i)}, d^{(i)})/r_1^{(i)})$, and the conclusion follows from the same argument given for Corollary 4.8. □

5 The tilt formula

Proposition 4.9 shows that, to determine whether a geometric decomposition of a hyperbolic manifold is Kojima's canonical one, we must lift the decomposition to Minkowski 4-space and then check convexity of all the angles at the 2-faces of the lifting. In this section we provide the explicit formula that allows to check convexity. This formula, already obtained by Ushijima [15] in more implicit terms, extends Weeks' tilt formula [16]. Our main contribution here is the computations of tilts in terms of moduli. The statements of Proposition 5.1, Remark 5.2, Theorem 5.4, and Propositions 5.5 and 5.6, with notation as in equation (11), may already be sufficient to proceed to the next section.

Tilts and convexity Let $\widehat{\Delta}$ be a tetrahedron in $\mathbb{M}^{3+1}$ that projects in Π^3 to a tetrahedron Δ with vertices outside $\mathbb{H}^3_{\mathrm{proj}}$ or on its boundary, and edges meeting $\partial\mathbb{H}^3_{\mathrm{proj}}$ or tangent to it. Let $\widehat{F}$ be a face of $\widehat{\Delta}$ with image F in Δ. Let H be the unique half-space in $\mathbb{H}^3_{\mathrm{proj}}$ such that $H \supset \Delta \cap \mathbb{H}^3_{\mathrm{proj}}$ and $\partial H \supset F \cap \mathbb{H}^3_{\mathrm{proj}}$. Let $m \in \mathcal{H}^3_+$ be the dual to H, and let $p \in \mathbb{M}^{3+1}$ be the unique vector such that $\langle p, x\rangle = -1$ for every $x \in \widehat{\Delta}$. We define the *tilt* of $\widehat{\Delta}$ relative to $\widehat{F}$ as the real number $\langle m, p\rangle$. The next result shows how tilts relate to convexity. For a proof see [15] or [16].

Proposition 5.1. *Let $\widehat{\Delta}$ and $\widehat{\Delta}'$ be tetrahedra in $\mathbb{M}^{3+1}$ sharing a 2-face $\widehat{F}$. Assume that $\widehat{\Delta} \cup \widehat{\Delta}'$ projects injectively to Π^3, and let t and t' be the tilts of $\widehat{\Delta}$ and $\widehat{\Delta}'$ relative to $\widehat{F}$. Then the angle formed by $\widehat{\Delta}$ and $\widehat{\Delta}'$ at $\widehat{F}$ is convex (respectively, strictly convex) if and only if $t + t' \leqslant 0$ (respectively, $t + t' < 0$).*

From moduli to tilts Our task is now to compute the tilts of the lifting of a partially truncated tetrahedron from the intrinsic geometry of the tetrahedron itself. Recall that the lifting of a non-ideal vertex u is uniquely determined by the requirement that it should belong to $\mathcal{H}^3_+$. However, when u is ideal, to get uniqueness we must choose a horosphere at u and lift u to the dual in L^3_+ to this horosphere. We begin by fixing some notation and recalling how horospheres are encoded [16].

Let Δ be an abstract partially truncated tetrahedron. Fix $\theta : \Delta^{(1)} \to [0, \pi)$ as in Theorem 2.2 and denote by $\Delta^{\theta,*}$ the corresponding geometric realization (up to isometry) of Δ^* in $\mathbb{H}^3_{\mathrm{proj}}$. Let Δ^θ be the associated tetrahedron with ideal and ultra-ideal vertices in Π^3.

Remark 5.2. If u is an ideal vertex of Δ^θ, the set of horospheres at u is parameterized by the positive reals, with a horosphere O corresponding to $r > 0$ if r is the radius of the smallest Euclidean disc on O containing $O \cap \Delta^\theta$.

So we fix a function $r : \mathcal{I} \to (0, \infty)$, where $\mathcal{I}$ is the set of ideal vertices of Δ, and we denote by $\mathcal{O}^{\theta,r}$ the associated family of horospheres at the ideal vertices of Δ^θ. Now $\mathcal{O}^{\theta,r}$ determines a unique lifting $\widehat{\Delta}^{\theta,r} \subset \mathbb{M}^{3+1}$ of Δ^θ, and $(\Delta^\theta, \widehat{\Delta}^{\theta,r})$ is well-defined up to isometry of pairs. Denoting by $u_1, \ldots, u_4$ the vertices of Δ and by F_i the face opposite to u_i, we can then define $t_i^{\theta,r}$ as the tilt of $\widehat{\Delta}^{\theta,r}$ relative to $\widehat{F}_i^{\theta,r}$. The tilt is unchanged under isometry, so $t_i^{\theta,r}$ is indeed well-defined.

To compute $t_i^{\theta,r}$ explicitly in terms of θ and r we must introduce certain positive real numbers $D_i^{\theta,r}$. To this end, we denote by H_i^θ the plane in $\mathbb{H}^3_{\mathrm{proj}}$ containing $F_i^{\theta,*}$. When u_i is non-ideal we denote by T_i^θ the truncation plane for $\Delta^{\theta,*}$ corresponding to u_i, and when u_i is ideal we denote by $O_i^{\theta,r}$ the horosphere at u_i determined by θ and r. Then we define $D_i^{\theta,r}$ as $D(H_i^\theta, T_i^\theta)$ or $D(H_i^\theta, O_i^{\theta,r})$ depending on the type of u_i, where D is the function defined as follows:

- If H and T are planes in $\mathbb{H}^3$ and $H \cap T = \emptyset$, we set $D(H, T) = \cosh d(H, T)$ where d is the usual distance in $\mathbb{H}^3$; if $H \cap T \neq \emptyset$ we set $D(H, T) = \cos \angle(H, T)$ where $\angle(H, T) \in (0, \pi/2]$ is the angle formed by H and T;

- If H is a plane in $\mathbb{H}^3$ and O is a horosphere not centred at a point of H, we set $D(H, O) = \exp \pm \ell$, where ℓ is the length of the unique geodesic arc that joins H to O and is orthogonal to both, with negative sign taken when $H \cap O \neq \emptyset$.

Remark 5.3. The choice of $\angle(H,T)$ in $(0,\pi/2]$ may look artificial at first sight, and indeed one could extend the definition to half-spaces rather than planes, and choose $\angle(H,T)$ in $(0,\pi)$. However it is easy to show that, in our situation, the dihedral angle at H and T that contains $\Delta^{\theta,*}$ is always the acute one, so the definition of D would remain the same.

The following result was proved in [15]:

Theorem 5.4. *Let θ and r determine the geometry and a family of horospheres at the ideal vertices of a partially truncated tetrahedron Δ with vertices $u_1,\ldots,u_4$. Let $\widehat{\Delta}^{\theta,r}$ be the tetrahedron in $\mathbb{M}^{3+1}$ determined by θ and r. Let $t_i^{\theta,r}$ be the tilt of $\widehat{\Delta}^{\theta,r}$ relative to the face opposite to u_i. Set $\theta_{ij}=\theta([u_i,u_j])$ and let $D_i^{\theta,r}$ be the number just introduced. Then:*

$$\begin{pmatrix} t_1^{\theta,r} \\ t_2^{\theta,r} \\ t_3^{\theta,r} \\ t_4^{\theta,r} \end{pmatrix} = \begin{pmatrix} 1 & -\cos\theta_{34} & -\cos\theta_{24} & -\cos\theta_{23} \\ -\cos\theta_{34} & 1 & -\cos\theta_{14} & -\cos\theta_{13} \\ -\cos\theta_{24} & -\cos\theta_{14} & 1 & -\cos\theta_{12} \\ -\cos\theta_{23} & -\cos\theta_{13} & -\cos\theta_{12} & 1 \end{pmatrix} \begin{pmatrix} 1/D_1^{\theta,r} \\ 1/D_2^{\theta,r} \\ 1/D_3^{\theta,r} \\ 1/D_4^{\theta,r} \end{pmatrix}.$$

To make the calculation of tilts explicit, we are left to compute the $D_i^{\theta,r}$'s. For ideal u_i we denote $r(u_i)$ by r_i, and the computation is easy:

Proposition 5.5. *If u_1 is ideal then:*

$$D_1^{\theta,r}=\frac{1}{2r_1}\cdot\frac{\sin\theta_{12}\cos\theta_{34}+\sin\theta_{13}\cos\theta_{24}+\sin\theta_{14}\cos\theta_{23}}{\sin\theta_{12}\sin\theta_{13}\sin\theta_{14}}.$$

Proof. We realize $\Delta^{\theta,*}$ in $\mathbb{H}^3_{\mathrm{half}}$ setting $u_1=\infty$ and we denote by C_i the trace at infinity of the plane that contains $F_i^{\theta,*}$. If a is the circumradius of the Euclidean triangle determined by C_2,C_3,C_4 and a' is the Euclidean radius of C_1, then it is readily shown that $D_1^{\theta,r}=a/(r_1\cdot a')$. Now the conclusion follows from the computation of a/a' starting from the moduli, which involves only tools of elementary Euclidean geometry, see Fig. 9. □

For non-ideal u_i we denote $D_i^{\theta,r}$ by D_i^{θ}, because it is independent of r. To compute it we need to introduce the following constant $g^\theta>0$:

$$\begin{aligned} g^\theta = -1 + & \sum\left\{\cos^2\theta(e):\ e\in\Delta^{(1)}\right\} \\ +2 & \sum\left\{\cos\theta(e')\cos\theta(e'')\cos\theta(e'''):\quad v\in\Delta^{(0)},\ v=e'\cap\ e''\cap e'''\right\} \\ +2 & \sum\left\{\frac{\prod\{\cos\theta(e):\ e\in\Delta^{(1)}\}}{\cos\theta(e')cos\theta(e'')}:\ \ \{e',e''\}\subset\Delta^{(1)},\ e'\cap e''=\emptyset\right\} \\ - & \sum\left\{\cos^2\theta(e')\cos^2\theta(e''):\ \ \{e',e''\}\subset\Delta^{(1)},\ e'\cap e''=\emptyset\right\}. \end{aligned} \tag{11}$$

Of course g^θ is well-defined. Moreover:

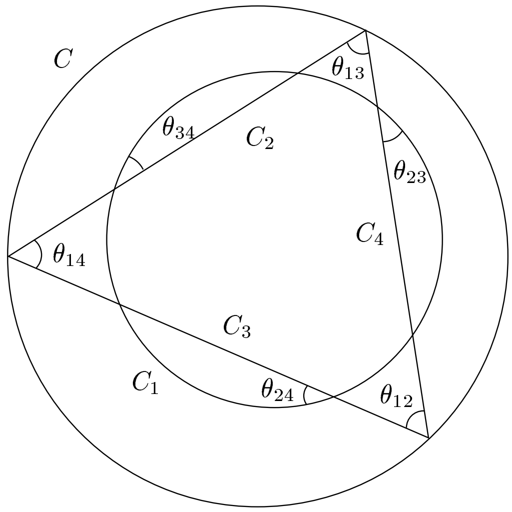

Figure 9: The ratio of the radii of C and C_1 is function of the θ_{ij}'s.

Proposition 5.6. *If u_i is non-ideal and d^θ is as in formula (3) of Section 2 then*

$$D_i^\theta = \sqrt{g^\theta / d^\theta(u_i)}.$$

The proof of this result will be divided in several lemmas. In the course of our argument we will need to use twice the following explicit formula for the hyperbolic distance in $\mathbb{H}^2_{\text{half}} = \{z \in \mathbb{C} : \Im(z) > 0\}$: for $x \in \mathbb{R}$, $\rho > 0$ and $0 < \alpha, \beta < \pi$ we have

$$d\big(x + \rho e^{i\alpha}, x + \rho e^{i\beta}\big) = \big|\log \tan(\alpha/2) - \log \tan(\beta/2)\big|. \tag{12}$$

Lemma 5.7. *Let $c, a > 0$ with $|c - a| < 1$. In $\mathbb{H}^2_{\text{half}}$ let γ_1 and γ_2 be the geodesics with ends at ± 1 and at $c \pm a$ respectively.*

1. *If γ_1 and γ_2 are disjoint then* $\cosh d(\gamma_1, \gamma_2) = (1 + a^2 - c^2) / 2a$.
2. *If γ_1 and γ_2 intersect at z and $\alpha = \angle(c - a, z, -1)$ is the angle they form then* $\cos \alpha = (1 + a^2 - c^2) / 2a$.

Proof. For point (1), let $p_1, p_2, \beta_1, \beta_2, x$ be as in Fig. 10. Since $\sin \beta_1 = 1/x$ and $\sin \beta_2 = a/(x - c)$, using (12) we easily get $d(\gamma_1, \gamma_2) = d(p_1, p_2) = \text{arccosh}\, x - \text{arccosh}\, ((x - c)/a)$. It follows that

$$\cosh d(\gamma_1, \gamma_2) = \left(x(x - c) - \sqrt{(x^2 - 1)((x - c)^2 - a^2)}\right) \Big/ a. \tag{13}$$

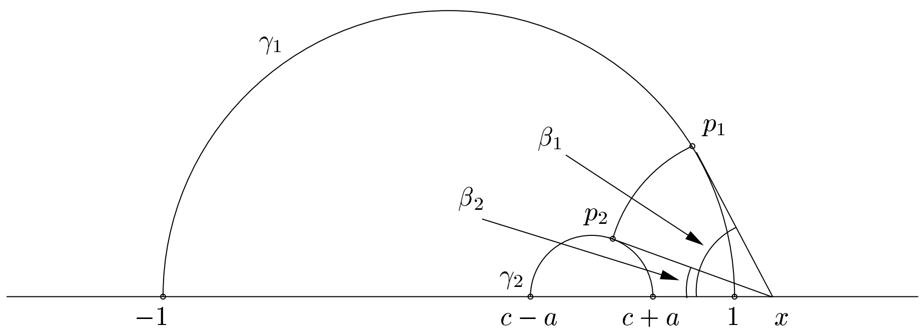

Figure 10: The geodesic arc from p_1 to p_2 is the shortest path between γ_1 and γ_2.

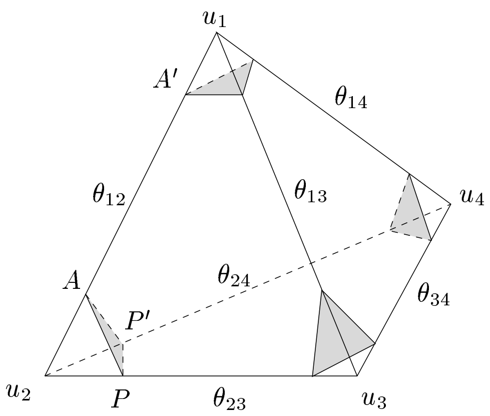

Figure 11: Notations for the proof of Proposition 5.6.

On the other hand, imposing that $|x-p_1|^2 = |x-p_2|^2$ we get that $x^2-1 = (x-c)^2-a^2$. Using this relation in the right-hand side of (13) we easily get the claimed equality. A very similar argument proves point (2). $\square$

Now assume that Δ has neither ideal vertices nor length-0 edges, *i.e.* that Δ^* is compact, and let A, A', P, P' be the points of Δ^* shown in Fig. 11. Let $\Delta^{\theta,*}$ be a geometric realization of Δ in $\mathbb{H}^3_{\rm half} = \mathbb{C} \times (0,\infty)$ such that $A = (0, \exp(-L^\theta([u_1,u_2])))$ and $A' = (0,1)$. Let the hemisphere H_1^θ containing $F_1^{\theta,*}$ have Euclidean radius R and centre $C = (z,0)$. Note that the truncation planes T_1^θ and T_2^θ relative to u_1 and u_2 are hemispheres centred at $(0,0)$ with Euclidean radii 1 and $\exp(-L^\theta([u_1,u_2]))$ respectively.

Proposition 5.8. $D_1^\theta = \frac{1}{R} \cdot \exp(-L^\theta([u_1,u_2])) \cdot \sinh(L^\theta([u_1,u_2]))$.

Proof. Since $H_1^\theta \perp T_2$, we have $|z|^2 = R^2 + \exp(-2L^\theta([u_1,u_2]))$. Lemma 5.7 with $c = |z|$ and $a = R$ yields $D_1^\theta = \big(1 - (\exp -2L^\theta([u_1,u_2]))\big) \,/\, 2R$. $\square$

Now let $\ell = d(A,P)$ and $\ell' = d(A',P')$. In the sequel we shall use the following equalities, which are readily deduced from [2, The Cosine Rule II, pag. 148]:

$$\tanh \ell = \frac{\sqrt{d^\theta(u_2)}}{\cos\theta_{12}\cdot\cos\theta_{23} + \cos\theta_{24}}, \qquad \tanh \ell' = \frac{\sqrt{d^\theta(u_2)}}{\cos\theta_{12}\cdot\cos\theta_{24} + \cos\theta_{23}}. \tag{14}$$

Proposition 5.9. $R^2 = \exp(-2L^\theta([u_1,u_2]))\cdot \sin^2\theta_{12} \,/\, d^\theta(u_2)$.

Proof. Using equality (12) it is easily seen that

$$\begin{aligned} P &= \exp(-L^\theta([u_1,u_2]))\,/\,\cosh\ell\cdot\big(\sinh\ell, 1\big),\\ P' &= \exp(-L^\theta([u_1,u_2]))\,/\,\cosh\ell'\cdot\big(\sinh\ell'\cos\theta_{12} + i\sinh\ell'\sin\theta_{12}, 1\big). \end{aligned}$$

Set $z = x+iy$. Since $0P \perp CP$ and $0P' \perp CP'$, we have

$$\begin{cases} x\cdot\tanh\ell = \exp(-L^\theta([u_1,u_2])),\\ x\cdot\tanh\ell'\cdot\cos\theta_{12} + y\cdot\tanh\ell'\cdot\sin\theta_{12} = \exp(-L^\theta([u_1,u_2])). \end{cases} \tag{15}$$

The desired equality is now readily proved by solving equations (15) with respect to x and y, using (14), and recalling that $R^2 = x^2+y^2-(\exp(-2L^\theta([u_1,u_2])))$. $\square$

We can now prove Proposition 5.6 for compact Δ^*. Equation (4) yields

$$\cosh L^\theta([u_1,u_2]) = c^\theta([u_1,u_2]) \Big/ \sqrt{d^\theta(u_1)d^\theta(u_2)},$$

where $c^\theta([u_1,u_2])$ is defined by equation (2). By Propositions 5.8 and 5.9 we deduce

$$\left(D_1^\theta\right)^2 = \frac{\cosh^2 L^\theta([u_1,u_2]) - 1}{R^2\cdot\exp 2L^\theta([u_1,u_2])} = \frac{\big(c^\theta([u_1,u_2])\big)^2 - d^\theta(u_1)d^\theta(u_2)}{d^\theta(u_1)\sin^2\theta_{12}}. \tag{16}$$

A long but straight-forward computation shows that the right-hand side of equation (16) is in fact equal to $g^\theta/d^\theta(u_1)$. This proves Proposition 5.6 when Δ is a truncated tetrahedron with no ideal vertices and no length-0 edges. In the general case we can approximate a geometric realization $\Delta^{\theta,*}$ of any partially truncated tetrahedron Δ with geometric realizations of compact truncated tetrahedra. Using Proposition 5.6 in the compact case and a standard continuity argument we then deduce that the proposition holds in general.

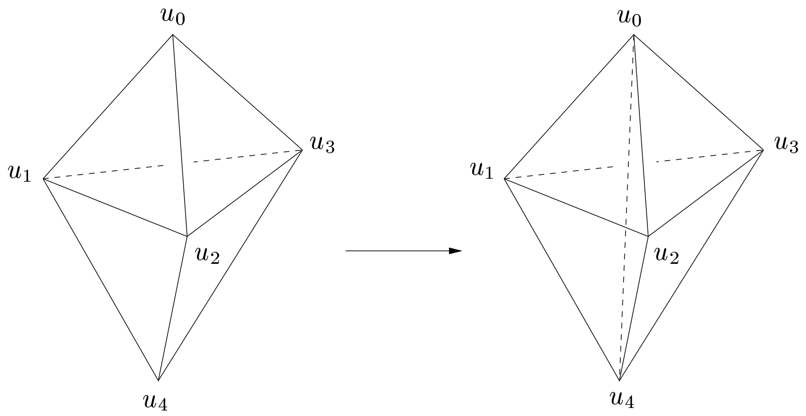

Figure 12: The two-to-three move.

6 Computing the canonical triangulation

In this section we show how to compute the canonical decomposition $\mathcal{K}_N^*$ of a hyperbolic 3-manifold N starting from an arbitrary geometric triangulation $\mathcal{T}^*$ of N. This is achieved by a step-by-step modification of $\mathcal{T}^*$ until a triangulation is reached whose lifting to $\mathbb{M}^{3+1}$ has only convex angles. According to Proposition 4.9, if all the angles are actually strictly convex then $\mathcal{K}_N^* = \mathcal{T}^*$, otherwise $\mathcal{K}_N^*$ is obtained from $\mathcal{T}^*$ by removing the 2-faces at which the lifting has flat angles. We warn the reader that, just as in [16], the process of transforming $\mathcal{T}^*$ into $\mathcal{K}_N^*$ may a priori get stuck at some point, so we are not entitled to call it an algorithm in a strict sense. On the other hand, in the next section we will show that if the process does not get stuck then it converges in finite time. The essential points of this section are the initial paragraph about topological and geometric moves, Theorem 6.6 and the outline of the algorithm described in the last two paragraphs.

Topological and geometric moves The fundamental move of the Matveev-Piergallini calculus for topological ideal triangulations is the two-to-three move, already mentioned in Section 1 and shown in Fig. 12. To fix notation, let us say that the move replaces two distinct tetrahedra T_{0123} and T_{1234} sharing a face F_{123} with three distinct tetrahedra T_{0124}, T_{0134}, and T_{0234} sharing an edge e_{04}. This move cannot always be performed in a geometric setting, but when it can we call the initial pair of tetrahedra an *admissible* one. More precisely:

Definition 6.1. A triple $(T_{0123}, T_{1234}, F_{123})$ consisting of two Euclidean tetrahedra and their common face, embedded in $\Pi^3 \cong \mathbb{E}^3$ as in Fig. 12-left, is called *admissible* if e_{04} meets the interior of F_{123}.

Recall now that we are considering a partially truncated triangulation $\mathcal{T}^*$ of a hyperbolic N. We denote by $\mathcal{T}$ the corresponding abstract triangulation, we fix the universal covering $\mathbb{H}^3_{\mathrm{proj}} \supset \widetilde{N} \to N$ and note that $\mathcal{T}$ determines a triangulation $\mathcal{T}_{\mathbb{P}}$ contained in Π^3. In the sequel we will often lift tetrahedra from $\mathcal{T}$ to $\mathcal{T}_{\mathbb{P}}$: the reader is invited to check that all our considerations are independent of the lifting chosen.

Definition 6.2. If F is a 2-face of $\mathcal{T}$ and the two tetrahedra Δ and Δ' incident to F are distinct, we call F *admissible* if the lifting to $\mathcal{T}_{\mathbb{P}} \subset \Pi^3 \cong \mathbb{E}^3$ of the triple (Δ, Δ', F) is a Euclidean admissible triple.

Remark 6.3. If F is an admissible 2-face of $\mathcal{T}$ then the two-to-three move that destroys F yields a new geometric partially truncated triangulation of N.

Turning to the inverse (three-to-two) move, we show that it is always geometric:

Lemma 6.4. *Assume in $\mathcal{T}$ there are precisely three distinct tetrahedra sharing a non-0-length edge. Lift the edge and the tetrahedra to $\mathcal{T}_{\mathbb{P}}$, with notation as in Fig. 12-right. Then the triple $(T_{0123}, T_{1234}, F_{123})$ is admissible, and the three-to-two move in $\mathcal{T}$ that destroys e_{04} gives rise to a new geometric triangulation of N.*

Proof. Since $T_{0124}, T_{0234}, T_{1234}$ are cyclically arranged around e_{04}, the line r_{04} through u_0 and u_4 meets the interior of F_{123}. To show that $r_{04} \cap F_{123}$ is actually a point of the interior of e_{04}, we must show that u_0 and u_4 cannot lie on opposite sides of the plane which contains F_{123}. If this were the case, using again the fact that $r_{04} \cap F_{123} \neq \emptyset$, we would deduce that $u_0 \in T_{1234}$ up to interchanging u_0 and u_4. From the Euclidean point of view, $u_0, \ldots, u_4$ lie outside the unit ball $\mathbb{H}^3_{\mathrm{proj}}$ or on its boundary, and all the edges e_{ij} meet the ball or are tangent to it. So $T_{1234} \setminus \mathbb{H}^3_{\mathrm{proj}}$ is the union of (at most) four regions W_i, where W_i is star-shaped with respect to u_i. Since u_0 belongs to one of the W_i's, the corresponding edge e_{0i} does not meet $\mathbb{H}^3_{\mathrm{proj}}$. A contradiction. $\square$

Effectiveness of moves To transform a geometric triangulation into Kojima's canonical decomposition, we will apply both the two-to-three and the three-to-two moves, trying to remove concave angles from the lifting to $\mathbb{M}^{3+1}$. The next result shows that, when we remove a concave angle, the new ones that we create are not concave. Note however that some of the "old" convex angles may become concave.

Lemma 6.5. *Let $(T_{0123}, T_{1234}, F_{123})$ be an admissible triple in Π^3, with notation as in Fig. 12-left. For $i = 0, \ldots, 4$ let $\widehat{u}_i$ be a lifting of u_i to $\mathbb{M}^{3+1}$. Assume that $\widehat{u}_0, \ldots, \widehat{u}_4$ are affinely independent in $\mathbb{M}^{3+1}$ and their convex hull X does not* *contain* 0. *Let* Y *be the set of almost-visible points of* X. *Then one and only one of the following possibilities occurs:*

- $Y = \widehat{T}_{0124} \cup \widehat{T}_{0134} \cup \widehat{T}_{0234}$, *the angles at* $\widehat{F}_{014}$, $\widehat{F}_{024}$, *and* $\widehat{F}_{034}$, *are strictly convex and the angle at* $\widehat{F}_{123}$ *is strictly concave.*

- $Y = \widehat{T}_{0123} \cup \widehat{T}_{1234}$,*the angle at* $\widehat{F}_{123}$ *is strictly convex and the angles at* $\widehat{F}_{014}$, $\widehat{F}_{024}$ *and* $\widehat{F}_{034}$ *are strictly concave.*

Proof. Let $Y' = \widehat{T}_{0124} \cup \widehat{T}_{0134} \cup \widehat{T}_{0234}$ and $Y'' = \widehat{T}_{0123} \cup \widehat{T}_{1234}$. Admissibility of $(T_{0123}, T_{1234}, F_{123})$ readily implies that either $Y = Y'$ or $Y = Y''$. More precisely, let $c = e_{04} \cap F_{123}$, $c' = \pi^{-1}(c) \cap \widehat{e}_{04}$, and $c'' = \pi^{-1}(c) \cap \widehat{F}_{123}$, and define $\lambda > 0$ so that $c'' = \lambda \cdot c'$. Affine independence of $\widehat{u}_0, \dots, \widehat{u}_4$ shows that $\lambda \neq 1$. Now $Y = Y'$ when $\lambda > 1$, and $Y = Y''$ when $\lambda < 1$.

On the other hand, choosing coordinates on $\mathbb{M}^{3+1}$ such that $\widehat{T}_{0123} \subset \{x_0 = 1\}$, we see that the angle at $\widehat{F}_{123}$ is convex precisely when $x_0(\widehat{u}_4) > 1$, namely when $\lambda < 1$. Similarly, with coordinates such that $\widehat{T}_{0124} \subset \{x_0 = 1\}$, the angle at $\widehat{F}_{024}$ is convex precisely when $x_0(\widehat{u}_3) > 1$, namely when $\lambda > 1$. This concludes the proof. □

Self-adjacent tetrahedra Now let $\mathcal{O}$ be a horospherical cross-section for N, and consider the corresponding lifting $\mathcal{T}_{\mathbb{M}}(\mathcal{O})$ of $\mathcal{T}_{\mathbb{P}}$ to $\mathbb{M}^{3+1}$. In $\mathcal{T}_{\mathbb{M}}(\mathcal{O})$ it is always possible to apply a two-to-three move to a 2-face. Similarly, we can always apply a three-to-two move to an edge with three neighbouring tetrahedra. The same moves may however be impossible in $\mathcal{T}$, when the involved tetrahedra are not distinct. The next result shows that in this case we actually do not need to worry about convexity of angles. In other words, when we are tempted to make a move (because of concavity), then we are guaranteed that the move is *topologically* possible. Recall however that a topological two-to-three move may not be geometric.

Theorem 6.6. *Let* $\mathcal{T}^*$ *be a geometric triangulation of a hyperbolic* N. *Take a universal cover* $\mathbb{H}^3_{\text{proj}} \supset \widetilde{N} \to N$, *let* $\mathcal{O}$ *be a horospherical cross-section for* N, *and let* $\mathcal{T}_{\mathbb{M}}(\mathcal{O})$ *be the associated lifting of* $\mathcal{T}$ *to* $\mathbb{M}^{3+1}$. *Then:*

1. *If* $\widehat{F}$ *is a* 2*-face of* $\mathcal{T}_{\mathbb{M}}(\mathcal{O})$ *and the two tetrahedra incident to* $\widehat{F}$ *have the same image in* $\mathcal{T}$, *then the angle at* $\widehat{F}$ *is strictly convex;*

2. *If* $\widehat{e}$ *is an edge of* $\mathcal{T}_{\mathbb{M}}(\mathcal{O})$ *with three incident tetrahedra* $\widehat{\Delta}_i$ *and three incident* 2*-faces* $\widehat{F}_i$, *for* $i = 1, 2, 3$, *and the images of the* $\widehat{\Delta}_i$*'s in* $\mathcal{T}$ *are not pairwise distinct, then the angles at the* $\widehat{F}_i$*'s are strictly convex.*

The proof of this result makes a crucial use of Proposition 2.7, announced but not shown in Section 2. So we show it now before proceeding.

Proof of 2.7. We use $\mathbb{H}^3_{\rm proj} \subset \Pi^3 \subset \mathbb{M}^{3+1}$ and take a tetrahedron Δ with vertices u_i outside $\mathbb{H}^3_{\rm proj} \cup \partial\mathbb{H}^3_{\rm proj}$ and edges e_{ij} that intersect $\mathbb{H}^3_{\rm proj}$ or are tangent to it. Let $\ell_{ij} \in [0,\infty)$ be the length of $e_{ij} \cap \Delta^*$. We must show that if Δ' is another such tetrahedron and $\ell'_{ij} = \ell_{ij}$ then there exists an isometry between Δ^* and Δ'^*.

Let $\widehat{u}_i$ be the only positive multiple of u_i that lies in $\mathcal{H}^3_+$. Since ℓ_{ij} is the distance between the truncation planes for Δ relative to u_i and u_j, and these planes bound the half-spaces dual to $\widehat{u}_i$ and $\widehat{u}_j$, Lemma 4.6 (1) shows that $\langle \widehat{u}_i, \widehat{u}_j \rangle = -\cosh \ell_{ij}$. This implies that $\langle \widehat{u}'_i, \widehat{u}'_j \rangle = \langle \widehat{u}_i, \widehat{u}_j \rangle$ for all i, j (including $i = j$). Now it is easy to see that $(\widehat{u}_i)_{i=1}^4$ and $(\widehat{u}'_i)_{i=1}^4$ are bases of $\mathbb{M}^{3+1}$. Then there exists an isometry φ of $\mathbb{M}^{3+1}$ such that $\varphi(\widehat{u}_i) = \widehat{u}'_i$ for all i, and the conclusion follows. $\square$

Proof of 6.6. For point (1), let $\widehat{\Delta}_1$ and $\widehat{\Delta}_2$ be the tetrahedra of $\mathcal{T}_{\mathbb{M}}(\mathcal{O})$ incident to $\widehat{F}$ and let Δ be their common image in $\mathcal{T}_{\mathbb{P}}$. Let $\theta : \Delta^{(1)} \to [0,\pi)$ give the dihedral angles of Δ and let $r : \mathcal{I} \to (0,\infty)$ be the map determined by $\mathcal{O}$ as explained in Remark 5.2, where $\mathcal{I}$ is the set of ideal vertices of Δ. Let $u_1, \ldots, u_4$ be the vertices of Δ and let F_i be the face opposite to u_i, with notation such that $(\widehat{\Delta}_i, \widehat{F})$ projects to (Δ, F_i) for $i = 1, 2$. According to Proposition 5.1, to prove strict convexity at $\widehat{F}$ we have to check that $t_1^{\theta,r} + t_2^{\theta,r} < 0$. To do so we will need to discuss various possibilities for the geometry of Δ. Recall first that the combinatorial data defining $\mathcal{T}$ determine a simplicial isomorphism $\varphi : F_1 \to F_2$, and φ induces an isometry $\varphi^* : F_1^{\theta,*} \to F_2^{\theta,*}$.

We consider first the case where $\theta_{34} = 0$. In this case Theorem 5.4 gives

$$\begin{aligned} t_1^{\theta,r} &= 1/D_1^{\theta,r} - 1/D_2^{\theta,r} - \cos\theta_{24}/D_3^{\theta,r} - \cos\theta_{23}/D_4^{\theta,r}, \\ t_2^{\theta,r} &= -1/D_1^{\theta,r} + 1/D_2^{\theta,r} - \cos\theta_{14}/D_3^{\theta,r} - \cos\theta_{13}/D_4^{\theta,r}. \end{aligned}$$

Now $\cos\theta_{24} + \cos\theta_{14} > 0$ and $\cos\theta_{23} + \cos\theta_{13} > 0$, so

$$t_1^{\theta,r} + t_2^{\theta,r} = -\Big((\cos\theta_{24} + \cos\theta_{14})/D_3^{\theta,r} + (\cos\theta_{23} + \cos\theta_{13})/D_4^{\theta,r} \Big) < 0.$$

The case $\theta_{34} = 0$ is settled, so we will assume henceforth that $\theta_{34} \neq 0$. This implies that $\varphi(u_2) \neq u_1$, otherwise the total dihedral angle in N along the image of e_{34} would reduce to θ_{34}, but $\theta_{34} < 2\pi$. Orientability of N then implies that, up to interchanging u_3 and u_4, we have

$$\varphi(u_2) = u_4, \ \varphi(u_3) = u_1, \ \varphi(u_4) = u_3. \tag{17}$$

These conditions easily imply that either all the u_i's are ideal or none of them is. Assume first they are all ideal. Then the dihedral angles along opposite edges are the same. We set $\alpha = \theta_{12} = \theta_{34}$, $\beta = \theta_{13} = \theta_{24}$, $\gamma = \theta_{14} = \theta_{23}$, and note further that $\alpha + \beta + \gamma = \pi$. Setting $r_i = r(u_i)$ and using Proposition 5.5 we see that $r_i = 1/D_i^{\theta,r}$. Recall now that the length of an edge of a Euclidean triangle is twice the circumradius times the sine of the opposite angle. Since φ^* matches the triangular cross-sections determined by the r_i's at the vertices of $\Delta^{\theta,*}$, we have:

$$r_2 \sin\alpha = r_4 \sin\beta, \quad r_3 \sin\beta = r_1 \sin\alpha, \quad r_4 \sin\gamma = r_3 \sin\gamma,$$

whence $r_2 = r_1$ and $r_4 = r_3 = r_1 \cdot \sin\alpha \,/\, \sin\beta$. Using Theorem 5.4 we then get

$$t_1^{\theta,r} = t_2^{\theta,r} = r_1 \cdot \big(1 - \cos\alpha - (\cos\beta + \cos\gamma) \cdot \sin\alpha \,/\, \sin\beta\big).$$

Relation $\alpha + \beta + \gamma = \pi$ now implies that $t_i^{\theta,r} = r_1 \cdot \sin\gamma \cdot (\cos\alpha - 1) \,/\, \sin\beta < 0$.

The only case left to settle to prove point (1) is when $\theta_{34} \neq 0$ and no u_i is ideal. Using (17) and the fact that φ^* is an isometry we see that

$$L^\theta(e_{23}) = L^\theta(e_{14}), \qquad L^\theta(e_{24}) = L^\theta(e_{34}) = L^\theta(e_{13}).$$

Proposition 2.7 just proved and these relations now imply that there exists an isometry of $\Delta^{\theta,*}$ onto itself that interchanges u_1 with u_2 and u_3 with u_4. Then, with notation as in Theorem 5.4, we have

$$D_1^\theta = D_2^\theta,\ D_3^\theta = D_4^\theta,\ \theta_{23} = \theta_{14},\ \theta_{13} = \theta_{24}.$$

Now we define $x = \cos\theta_{24} = \cos\theta_{13},\ y = \cos\theta_{23} = \cos\theta_{14},\ z = \cos\theta_{34},\ w = \cos\theta_{12}$. From Theorem 5.4 and Proposition 5.6 we deduce that

$$t_1^\theta = t_2^\theta = \Big((1-z)\sqrt{d^\theta(u_1)} - (x+y)\sqrt{d^\theta(u_3)}\Big) \Big/ \sqrt{g^\theta}. \tag{18}$$

We claim that the following formula holds:

$$w = (x + zy)\sqrt{(1-x^2)\,/\,(1-z^2)} - xy. \tag{19}$$

We first show that (19) implies $t_i^\theta < 0$. Later we will establish (19). Note that $d^\theta(u_1) = x^2 + y^2 + w^2 + 2xyw - 1$ and $d^\theta(u_3) = x^2 + y^2 + z^2 + 2xyz - 1$. Using (19) we then get $d^\theta(u_1) = d^\theta(u_3) \cdot \frac{1-x^2}{1-z^2}$. By equation (18) it follows that

$$t_i^\theta = \sqrt{d^\theta(u_3)} \cdot \Big((1-z)\sqrt{(1-x^2)\,/\,(1-z^2)} - (x+y)\Big) \Big/ \sqrt{g^\theta}.$$

Since $\theta_{14} + \theta_{24} + \theta_{34} < \pi$ we deduce that

$$x + y > 1 - z \geqslant 0. \tag{20}$$

Then $t_i^\theta < 0 \Leftrightarrow x + y > (1-z)\sqrt{\frac{1-x^2}{1-z^2}} \Leftrightarrow (x+y)^2 > \Big((1-z)\sqrt{\frac{1-x^2}{1-z^2}}\Big)^2$. After some computations we deduce that

$$t_i^\theta < 0 \Leftrightarrow (x^2 + y^2 + z^2 + 2xyz - 1) + (x+y)^2 + (1-z)(z-y)^2 > 0.$$

Since $x^2 + y^2 + z^2 + 2xyz - 1 = d^\theta(u_3) > 0$, it is now sufficient to show that

$$(x + y)^2 + (1 - z)(z - y^2) \geqslant 0,$$

and this inequality follows quite easily from (20).

We are left to establish (19). To do so we distinguish two cases, according to whether θ_{14} vanishes or not. We first assume $\theta_{14} \neq 0$. Let T_2 and T_4 be the truncation triangles relative to u_2 and u_4 respectively and let ℓ, ℓ' be the boundary edges of $\Delta^{\theta,*}$ defined by $\ell = F_1^{\theta,*} \cap T_2$, $\ell' = F_2^{\theta,*} \cap T_4$. Now $\theta_{34} \neq 0$ and $\theta_{14} \neq 0$, so ℓ' has finite length. Moreover, we know that φ^* identifies ℓ to ℓ'. Equation (19) is now obtained by equaling the lengths of ℓ and ℓ' via Proposition 2.4. We turn at last to the case when $\theta_{14} = 0$. Imposing $L^\theta(e_{24}) = L^\theta(e_{34})$ via Proposition 2.6 we get after some computations $(x + w)^2(1 - z^2) = (x + z)^2(1 - x^2)$. Since $y = 1$, this relation is in fact equivalent to equation (19), and we are done.

Now we show point (2). As a consequence of (1), we have strict convexity along at least one of the $\widehat{F}_i$'s. Lemma 6.5 then implies strict convexity at all the $\widehat{F}_i$'s. $\square$

Outline of the algorithm The input of our process is a geometric triangulation $\mathcal{T}^*$ of a hyperbolic 3-manifold N. More precisely, we start with a topological partially truncated triangulation $\mathcal{T}$ of N in the sense of Proposition 1.15, and a solution of the system of consistency and completeness equations for $\mathcal{T}$, as discussed in Section 2. Then we perform the following steps:

1. We choose a horospherical cross-section $\mathcal{O}$ as in Proposition 4.7;

2. We pick a 2-face F of $\mathcal{T}$ such that the two tetrahedra of $\mathcal{T}$ incident to F are distinct. We lift F and its two incident tetrahedra to $\mathcal{T}_{\mathbb{M}}(\mathcal{O})$. Using Proposition 5.1, we check whether the angle at the lifted face is strictly concave. If it is, we move to step 3. If it is not, we move to another 2-face. If all faces are visited and no concave angle is found, $\mathcal{T}^*$ is the output (because it is $\mathcal{K}_N^*$ or a subdivision of it, by Proposition 4.9 and Theorem 6.6);

3. If F is admissible, we change $\mathcal{T}$ by performing the geometric two-to-three move that kills F, and we go back to step 2. If F is non-admissible, we check whether one of the non-0-length edges of F is shared by precisely three tetrahedra of $\mathcal{T}$. If it is, we change $\mathcal{T}$ by applying the geometric three-to-two move that kills this edge, and we go back to step 2. If it is not, we do not change $\mathcal{T}$ but we go back to step 2 moving to a different concave face. If all concave faces are visited and no move can be applied to any of them, we give up.

In the next section we will show that the process, if it does not get stuck during step 3, outputs the canonical decomposition in finite time. Steps 2 and 3 are of course directly implementable, whereas step 1 requires a careful discussion, to which the rest of the present section is devoted.

Algorithmic choice of horospherical cross-sections According to Proposition 4.11, to determine a horospherical cross-section as in Proposition 4.7 we must find for each cusp a realization of $\widetilde{N}$ in $\mathbb{H}^3_{\text{half}}$ so that the cusp is generated by ∞, and compute the corresponding d, r_1, r_2 of Proposition 4.10. We recall that the datum to use is a geometric triangulation $\mathcal{T}$ of N.

Let us concentrate on a cusp C and fix a tetrahedron $\Delta_0^* \in \mathcal{T}^*$ with a certain ideal vertex v_0 asymptotic to C. We take a realization $\tilde{\Delta}_0^*$ in $\mathbb{H}^3_{\text{half}}$ such that v_0 gets identified to ∞. Here and in the sequel the realizations we consider are of course all compatible with the geometric structure given on the tetrahedra. Choosing a horospherical cross-section at C now amounts to choosing a positive real number, namely the height at which the lifted cross-section should intersect $\tilde{\Delta}_0$. In the course of our argument, starting from $\tilde{\Delta}_0^*$, we will be successively gluing new tetrahedra to free faces of tetrahedra we already have, as dictated by the combinatorics and the geometry of $\mathcal{T}^*$. We warn the reader that it is not possible to predict *a priori* how many different copies of each tetrahedron of $\mathcal{T}^*$ will need to be glued, but the process is guaranteed to be finite anyway, as we will carefully explain.

Step 1.A. *We take one copy of each tetrahedron Δ^* of $\mathcal{T}^*$ for each vertex v of Δ^* asymptotic to C, and, starting from $\tilde{\Delta}_0^*$, we do gluings along free vertical faces until each (Δ^*, v) has been realized once in $\mathbb{H}^3_{\text{half}}$ with $v = \infty$.*

After Step 1.A we have a certain finite family $\mathcal{F}_1$ of partially truncated tetrahedra in $\mathbb{H}^3_{\text{half}}$, all having ∞ as a vertex, and we can compute the following:

- $\rho = \max\{\rho(\tilde{\Delta}^*) :\ \tilde{\Delta}^* \in \mathcal{F}_1\}$, where $\rho(\tilde{\Delta}^*)$ is the Euclidean radius of the half-sphere that contains the face of $\tilde{\Delta}^*$ opposite to ∞;
- $r = \max \bigcup \{r(\tilde{\Delta}^*) :\ \tilde{\Delta}^* \in \mathcal{F}_1\}$, where $r(\tilde{\Delta}^*)$ is the set of Euclidean radii of the half-spheres that contain the truncation triangles of $\tilde{\Delta}^*$. We define r to be $-\infty$ if all the tetrahedra of $\mathcal{F}_1$ are ideal, and we note for later purpose that the definition of $r(\tilde{\Delta}^*)$ makes sense also if $\tilde{\Delta}^*$ does not have ∞ as a vertex;
- The intersection Ω of the horizontal plane at height $z = \max\{\rho, r\}$ with the union of the tetrahedra in $\mathcal{F}_1$;
- The first number we need to determine, *i.e.* the diameter d of Ω with respect to the ordinary Euclidean metric on $\mathbb{C} \times \{z\}$.

Step 1.B. *Starting from $\mathcal{F}_1$, we perform gluings along free non-vertical faces, adding new truncated tetrahedra, until we get a family $\mathcal{F}_2$ such that $\bigcup\{r(\tilde{\Delta}^*) :\ \tilde{\Delta}^* \in \mathcal{F}_2\}$ contains at least two distinct values $r_1' > r_2'$.*

The way to realize Step 1.B algorithmically is as follows. We list the free non-vertical faces of $\mathcal{F}_1$, we perform the gluings along these faces getting a family $\mathcal{F}_1'$, and we check whether $\mathcal{F}_1'$ already works. If it does not, we proceed similarly with $\mathcal{F}_1'$ instead of $\mathcal{F}_1$, until the desired $\mathcal{F}_2$ is reached. Of course this procedure only has to be iterated a finite number of times, even if the number of iterations is hard to predict *a priori*.

Step 1.C. *Starting from $\mathcal{F}_2$, we perform gluings along free non-vertical faces, adding new truncated tetrahedra, until we get a family $\mathcal{F}_3$ such that any further tetrahedron glued to $\mathcal{F}_3$ along a non-vertical face would lie entirely outside $\Omega \times [r_2', \infty)$.*

Of course this step is also a finite one, even if its length is not easily predictable. Note also that Ω has finite diameter d, so we could replace Ω by an easier set, like a disc or a square. The choice of $\mathcal{F}_3$ guarantees that its union contains $\widetilde{N} \cap (\Omega \times [r_2', \infty))$, so the two other constants r_1 and r_2 we need to determine are now the first and second largest elements of $\bigcup \{r(\tilde{\Delta}^*) : \ \tilde{\Delta}^* \in \mathcal{F}_3\}$.

7 Finiteness

This section is entirely devoted to the proof that the algorithm to transform a geometric triangulation into Kojima's canonical decomposition, if it does not get stuck, converges in finite time. This fact was already announced above and is accurately stated as follows:

Theorem 7.1. *Let N be hyperbolic with non-empty boundary, let $\mathcal{O}$ be a horospherical cross-section for N as in Proposition 4.7, and let $\mathcal{T}$ be a geometric triangulation of N. Then there exists an integer $\nu = \nu(N, \mathcal{O}, \mathcal{T})$ such that the following holds: Assume $\{\mathcal{T}^i\}_{i=0}^{j}$ is a sequence of geometric triangulations of N starting at $\mathcal{T}^0 = \mathcal{T}$, and for all $i = 0, \ldots, j-1$ we have:*

- *there is a 2-face $\widehat{F}^i$ of $\mathcal{T}^i_{\mathbb{M}}(\mathcal{O})$ along which $\mathcal{T}^i_{\mathbb{M}}(\mathcal{O})$ has a strictly concave angle;*
- *$\mathcal{T}^{i+1}$ is obtained from $\mathcal{T}^i$ by a two-to-three or a three-to-two move killing $\widehat{F}^i$.*

Then $j \leqslant \nu$.

For the proof we fix as above the universal cover $\mathbb{H}^3 \supset \widetilde{N} \to N$ with deck transformation group $\Gamma < \mathrm{Isom}(\mathbb{H}^3)$, and we denote by $\widetilde{\mathcal{O}}$ the lifting of $\mathcal{O}$. We start with a series of lemmas, the first of which is taken from [8].

Lemma 7.2. *A point $q \in \widetilde{N}_\infty$ generates an annular cusp of N if and only if it belongs to the circle at infinity of two distinct components of $\partial \widetilde{N}$.*

Lemma 7.3. *Let S_0 be a component of $\partial\widetilde{N}$ with stabilizer Γ_0 in Γ. If $c > 0$ and*

$$A(S_0, c) = \{S \subset \partial\widetilde{N} : \ S \text{ is a component of } \partial\widetilde{N}, \ d(S, S_0) \leqslant c\},$$

then Γ_0 leaves $A(S_0, c)$ invariant and $\#(A(S_0, c)/\Gamma_0) < \infty$.

Proof. The first assertion is obvious. For $S \in A(S_0, c)$ we have either $d(S, S_0) > 0$ or $d(S, S_0) = 0$. Correspondingly we have a splitting $A(S_0, c) = A^+(S_0, c) \sqcup A^0(S_0)$, which of course is Γ_0-equivariant. Recall now from [8] that $S_0/\Gamma_0 = S_0/\Gamma$ is a complete finite-area hyperbolic surface. Lemma 7.2 readily implies that there is a bijection between $A^0(S_0)/\Gamma_0$ and the set of cusps of S_0/Γ_0, so $A^0(S_0)/\Gamma_0$ is finite.

We are left to show that $A^+(S_0, c)/\Gamma_0$ is finite. To this end note first that $\widetilde{N}/\Gamma_0$ is complete hyperbolic with geodesic boundary (but probably infinite volume), and its boundary components constitute a locally finite family. Now let q be a point of ∂S_0 that generates a cusp of S_0/Γ_0, *i.e.* a point in $\partial S_0 \subset \widetilde{N}_\infty$ that generates an annular cusp of N. Using Lemma 7.2 and realizing $\widetilde{N}$ in $\mathbb{H}^3_{\text{half}}$ with $q = \infty$, it is easily proved that there exists a horoball B_q centred at q such that, if S is a component of $\partial\widetilde{N}$ and $d(S, S_0) > 0$, then $d(S, S_0) = d(S, S_0 \setminus O_q)$. Repeating this argument for all the finitely many cusps of S_0/Γ_0 we deduce that for some $\varepsilon > 0$ our quotient $A^+(S_0, c)/\Gamma_0$ naturally corresponds to the set of boundary components of $\widetilde{N}/\Gamma_0$ whose distance from the ε-thick part of S_0/Γ_0 is positive and bounded by c. Compactness of the ε-thick part of S_0/Γ_0 and local finiteness of the components of $\partial(\widetilde{N}/\Gamma_0)$ then imply the conclusion. $\square$

Lemma 7.4. *Let O_0 be a component of $\widetilde{\mathcal{O}}$ with stabilizer Γ_0 in Γ. If $c > 0$ and*

$$\begin{aligned} A^{(1)}(O_0, c) &= \{O \subset \widetilde{\mathcal{O}} : \ O \text{ is a component of } \widetilde{\mathcal{O}}, \ d(O, O_0) \leqslant c\}, \\ A^{(2)}(O_0, c) &= \{S \subset \partial\widetilde{N} : \ S \text{ is a component of } \partial\widetilde{N}, \ d(S, O_0) \leqslant c\}, \end{aligned}$$

then Γ_0 leaves $A^{(j)}(O_0, c)$ invariant and $\#(A^{(j)}(S_0, c)/\Gamma_0) < \infty$ for both $j = 1, 2$.

Proof. Realize $\widetilde{N}$ in $\mathbb{H}^3_{\text{half}} = \mathbb{C} \times (0, \infty)$ so that O_0 is centred at ∞. Now the components of $\widetilde{O}$ are Euclidean spheres, and $A^{(1)}(O_0, c)$ consists of those whose radius is bounded from below by a certain constant. Compactness of $\mathbb{C}/\Gamma_0$ easily implies finiteness of $A^{(j)}(S_0, c)/\Gamma_0$ for $j = 1$. A very similar argument is employed for $j = 2$. $\square$

Now we set $\mathcal{X}_{\mathcal{T}}(\mathcal{O}) = \{\alpha \cdot x : \ x \in \mathcal{T}_{\mathbb{M}}(\mathcal{O}), \ \alpha \geqslant 1\}$ and for real c we define $\mathcal{L}_c = \{v \in \mathbb{M}^{3+1} : \ x_0(v) > 0, \ \langle v, v\rangle < c\}$.

Lemma 7.5. *There exists $c < 0$ such that $\mathcal{X}_{\mathcal{T}}(\mathcal{O}) \supset \mathcal{L}_c$.*

Proof. Since Γ acts isometrically on $\mathcal{T}_{\mathbb{M}}(\mathcal{O})$, the Lorentzian norm induces a continuous map $\mathcal{T}_{\mathbb{M}}(\mathcal{O}) \to \mathbb{R}$. The domain of this map is homeomorphic to a compactification of N, so the map has a minimum c, which of course is negative. Knowing that the projection of $\mathcal{T}_{\mathbb{M}}(\mathcal{O})$ to Π^3 contains the unit ball $\mathbb{H}^3_{\text{proj}}$ we easily deduce the conclusion. $\square$

Proposition 7.6. *There exist only a finite number of* Γ*-inequivalent segments with ends in* $\mathcal{B} \cup \widetilde{Q}(\mathcal{O})$ *whose midpoint does not belong to* $\mathcal{X}_{\mathcal{T}}(\mathcal{O})$.

Proof. Let $p_1, p_2 \in \mathcal{B} \cup \widetilde{Q}(\mathcal{O})$ be distinct and assume that the midpoint of $[p_1, p_2]$ does not belong to $\mathcal{X}_{\mathcal{T}}(\mathcal{O})$. Choosing c as in Lemma 7.5 we deduce that the Lorentzian norm of $(p_1 + p_2)/2$ is at least c. We will now consider three cases depending on whether p_1 and p_2 belong to $\mathcal{B}$ or to $\widetilde{Q}(\mathcal{O})$. In all three cases we will show that there are finitely many choices for $\{p_1, p_2\}$ up to the action of Γ.
Case 1: $p_1, p_2 \in \mathcal{B}$. Then $c \leqslant \big\langle (p_1 + p_2)/2, (p_1 + p_2)/2 \big\rangle = \big(1 + \langle p_1, p_2 \rangle\big)/2$. If S_i is the component of $\partial\widetilde{N}$ dual to p_i, using Lemma 4.6 (1) we deduce that $\cosh d(S_1, S_2) \leqslant 1 - 2c$. Finiteness of the number of components of ∂N and Lemma 7.3 then imply the desired finiteness.
Case 2: $p_1, p_2 \in \widetilde{Q}(\mathcal{O})$. Then $c \leqslant \big\langle (p_1 + p_2)/2, (p_1 + p_2)/2 \big\rangle = \langle p_1, p_2 \rangle/2$. If O_i is the horosphere dual to p_i, we deduce that $\exp d(O_1, O_2) \leqslant -c$ by Lemma 4.6 (4). We then get the desired conclusion using finiteness of the number of toric cusps of N and Lemma 7.4 with $j = 1$.
Case 3: $p_1 \in \widetilde{Q}(\mathcal{O}), p_2 \in \mathcal{B}$. Then $c \leqslant \big\langle (p_1 + p_2)/2, (p_1 + p_2)/2 \big\rangle = \big(1 + 2\langle p_1, p_2 \rangle\big)/4$. If O is the horosphere dual to p_1 and S is the component of $\partial\widetilde{N}$ dual to p_2, we deduce that $2 \cdot \exp d(O, S) \leqslant 1 - 4c$ by Lemma 4.6 (3). We conclude using again finiteness of the number of toric cusps of N and Lemma 7.4 with $j = 2$. $\square$

Proof of 7.1. Lemma 6.5 implies that $\mathcal{X}_{\mathcal{T}^i}(\mathcal{O}) \subset \mathcal{X}_{\mathcal{T}^{i+1}}(\mathcal{O})$ for all i. Moreover $\mathcal{X}_{\mathcal{T}^{i+1}}(\mathcal{O})$ contains at least an edge with endpoints in $\mathcal{B} \cup \widetilde{Q}(\mathcal{O})$ whose midpoint does not belong to $\mathcal{X}_{\mathcal{T}^i}(\mathcal{O})$. Then we achieve the desired property by defining ν as the number of Γ-inequivalent edges with endpoints in $\mathcal{B} \cup \widetilde{Q}(\mathcal{O})$ whose midpoints do not belong to $\mathcal{X}_{\mathcal{T}}(\mathcal{O})$. $\square$

References

[1] S. Baseilhac, R. Benedetti, *Quantum hyperbolic state sum invariants of 3-manifolds*, `math.GT/0101234`.

[2] A. F. Beardon, "The geometry of discrete groups", Graduate Texts in Mathematics, Vol. 91, Springer-Verlag, New York, 1995.

[3] R. Benedetti, C. Petronio, "Lectures in Hyperbolic Geometry", Universitext, Springer-Verlag, Berlin, 1992.

[4] P. J. Callahan, M. V. Hildebrandt, J. R. Weeks, *A census of cusped hyperbolic* 3*-manifolds. With microfiche supplement*, Math. Comp. **68** (1999), 321–332.

[5] D. B. A. Epstein, R. C. Penner, *Euclidean decompositions of noncompact hyperbolic manifolds*, J. Differential Geom. **27** (1988), 67–80.

[6] M. Fujii, *Hyperbolic* 3*-manifolds with totally geodesic boundary*, Osaka J. Math. **27** (1990), 539–553.

[7] M. Fujii, *Hyperbolic* 3*-manifolds with totally geodesic boundary which are decomposed into hyperbolic truncated tetrahedra*, Tokyo J. Math. **13** (1990), 353–373.

[8] S. Kojima, *Polyhedral decomposition of hyperbolic manifolds with boundary*, Proc. Work. Pure Math. **10** (1990), 37–57.

[9] S. Kojima, *Polyhedral decomposition of hyperbolic* 3*-manifolds with totally geodesic boundary*, In: "Aspects of low-dimensional manifolds, Kinokuniya, Tokyo", Adv. Stud. Pure Math. **20** (1992), 93–112.

[10] S. V. Matveev, *Transformations of special spines, and the Zeeman conjecture*, Math. USSR-Izv. **31** (1988), 423–434.

[11] R. Piergallini, *Standard moves for standard polyhedra and spines*, In: "Third National Conference on Topology, Trieste, 1986", Rend. Circ. Mat. Palermo (2) Suppl. **18** (1988), 391–414.

[12] W. P. Thurston, "The Geometry and Topology of 3-manifolds", mimeographed notes, Princeton, 1979.

[13] W. P. Thurston, *Three-dimensional manifolds, Kleinian groups and hyperbolic geometry*, Bull. Amer. Math. Soc. (N.S.) **6** (1982), 357–381.

[14] V. G. Turaev, O.Ya. Viro, *State sum invariants of* 3*-manifolds and quantum* $6j$*-symbols*, Topology **31** (1992), 865–902.

[15] A. Ushijima, *A unified viewpoint about geometric objects in hyperbolic space and the generalized tilt formula*, In: "Hyperbolic spaces and related topics, II, Kyoto, 1999", Sūrikaisekikenkyūsho Kōkyūroku **1163** (2000), 85–98.

[16] J. R. Weeks, *Convex hulls and isometries of cusped hyperbolic 3-manifolds*, Topology Appl. **52** (1993), 127–149.

[17] J. Weeks, *SnapPea*, The hyperbolic structures computer program, available from `www.northnet.org/weeks`.

Scuola Normale Superiore
Piazza dei Cavalieri, 7
56126 Pisa, Italy
frigerio@sns.it

Dipartimento di Matematica Applicata
Università di Pisa
Via Bonanno Pisano, 25/B
56126 Pisa, Italy
petronio@dma.unipi.it