Triangulations of 3-manifolds, hyperbolic relative handlebodies, and Dehn filling

François Costantino Roberto Frigerio
Bruno Martelli Carlo Petronio

March 9, 2004

Abstract

We establish a bijective correspondence between the set $\mathcal{T}_n$ of 3-dimensional triangulations with n tetrahedra and a certain class $\mathcal{H}_n$ of relative handlebodies (*i.e.* handlebodies with boundary loops, as defined by Johannson) of genus $n+1$.

We show that the manifolds in $\mathcal{H}_n$ are hyperbolic (with geodesic boundary, and cusps corresponding to the loops), have least possible volume, and simplest boundary loops.

Mirroring the elements of $\mathcal{H}_n$ in their geodesic boundary we obtain a set $\mathcal{D}_n$ of cusped hyperbolic manifolds, previously considered by D. Thurston and the first named author. We show that also $\mathcal{D}_n$ corresponds bijectively to $\mathcal{T}_n$, and we study some Dehn fillings of the manifolds in $\mathcal{D}_n$. As consequences of our constructions, we also show that:

- A triangulation of a 3-manifold is uniquely determined up to isotopy by its 1-skeleton;
- If a 3-manifold M has an ideal triangulation with edges of valence at least 6, then M is hyperbolic and the edges are homotopically non-trivial, whence homotopic to geodesics;
- Every finite group G is the isometry group of a closed hyperbolic 3-manifold with volume less than const $\times\ |G|^9$.

MSC (2000): 57M50 (primary), 57M20, 57M25 (secondary).

Introduction

Let us extend the traditional notion of triangulation of a 3-manifold by calling *triangulation* any combinatorial pattern of face-pairings between a finite number of tetrahedra. So a triangulation may actually not define a manifold. We also define a *hyperbolic* manifold to be a complete finite-volume hyperbolic 3-manifold with (possibly empty) geodesic boundary.

In this paper we analyze the class $\mathcal{T}$ of triangulations in geometric terms, showing that it can be identified to a set $\mathcal{H}$ of hyperbolic 3-manifolds, and to the set $\mathcal{D}$ of their "doubles," already considered in [4] by D. Thurston and the first named author. The correspondences we construct shade new light on each of the objects we consider and allow us to prove several topological and geometric facts.

In the rest of the present introduction we state our main results in a roughly decreasing order of relevance. Proofs will be given in Sections 1 to 4, in a different order.

Relative handlebodies Following Johannson [12], a *relative handlebody* (H, Γ) is a (possibly non-orientable) handlebody H with a finite system Γ of disjoint loops on ∂H. A *hyperbolic structure* on (H, Γ) is a finite-volume complete hyperbolic structure with totally geodesic boundary on $H \setminus \Gamma$, so each loop in Γ geometrically corresponds to a cusp based on a strip (an annulus or a Möbius band).

We call *pant-meridinal complexity* of Γ the minimum $c(\Gamma)$ of $|\Gamma \cap \partial D|$ where D is a system of disjoint properly embedded discs in H such that ∂D cuts ∂H into a union of pairs of pants.

From triangulations to relative handlebodies If $T \in \mathcal{T}$ we now define a relative handlebody $N(T)$ as follows. First, we glue together the tetrahedra of T according to the face-pairings, getting the support $|T|$ of T. Then, we remove from $|T|$ an open neighbourhood of the vertices, getting a space $M(T)$. Next, we remove from $M(T)$ an open neighbourhood W of the edges, which gives a handlebody H. Last, we note that ∂W consists of strips, we define Γ as the system of cores of ∂W, and we set $N(T) = (H, \Gamma)$. We often denote by $N(T)$ also the non-compact manifold $H \setminus \Gamma$, the context making clear which definition of $N(T)$ we are referring to. An example of $N(T)$ is shown in Fig. 1.

In Section 1 we give $N(T)$ a hyperbolic structure, by taking a regular hyperbolic ideal octahedron for each tetrahedron of T, and gluing these octahedra appropriately. We then establish the following result. We denote by $\mathcal{T}_n$ the set of triangulations involving n tetrahedra.

Theorem 0.1. *The map $T \mapsto N(T)$ just defined gives a bijection between $\mathcal{T}_n$ and a set $\mathcal{H}_n$ of hyperbolic relative handlebodies of genus $n+1$. Among*

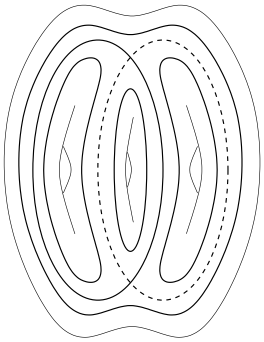

If T is the triangulation of S^3 obtained by gluing two tetrahedra along the identity of their boundary, the corresponding genus-3 relative handlebody $N(T) = (H, \Gamma)$ is as shown. Its geometric realization, which has 6 annular cusps and decomposes into two regular ideal octahedra, is also the building block of the elements of $\mathcal{D}$ (defined below).

Figure 1: An example of $N(T)$.

the hyperbolic (H, Γ)'s of genus $n+1$, the set $\mathcal{H}_n$ consists of:

- *those having minimal complexity $c(\Gamma)$, equal to $10 \cdot n$;*
- *those having minimal volume, equal to $n \cdot v_O$, where $v_O \approx 3.66386$ is the volume of a hyperbolic regular ideal octahedron.*

The map $N : \mathcal{T} \to \mathcal{H} := \bigcup_{n=1}^{\infty} \mathcal{H}_n$ is a bijection because the decomposition of $N(T)$ into octahedra is Kojima's canonical one [14]. A consequence of this bijection is the following result, which appears to be new and to have independent interest. If Y is a compact 3-manifold, we call T a *triangulation* of Y if a space homeomorphic to Y can be obtained from $|T|$ by removing open stars of some of the vertices. Therefore we have an ideal vertex for each component (if any) of ∂Y, and possibly some internal vertices. When we remove stars of all the vertices, we say that T is an *ideal* triangulation of Y, and Y coincides with the space $M(T)$ already defined above.

Proposition 0.2. *Let T_0 and T_1 be triangulations of the same compact 3-manifold Y. Assume that the 1-skeleta of T_0 and T_1 are the same graph contained in Y. Then T_0 and T_1 are isotopic relatively to the 1-skeleton.*

Tangles We now call *tangle* in a compact 3-manifold with boundary a finite union of disjoint properly embedded arcs. As above, the complement of an open tubular neighbourhood W of a tangle is a manifold with boundary loops given by the cores of ∂W. Since every tangle is contained in the 1-skeleton of an ideal triangulation [3], we have the following result.

Corollary 0.3. *Every tangle in every compact 3-manifold is contained in a tangle whose complement lies in $\mathcal{H}$ (in particular, it is hyperbolic).*

Mirroring If (H, Γ) is a relative handlebody we now define its "double" $D(H, \Gamma)$ as follows. We take the orientation covering $\widetilde{H}$ of H with deck involution τ (so $H = \widetilde{H}/_{\tau}$), we consider the pre-image $\widetilde{\Gamma}$ of Γ in $\widetilde{H}$, and we define $D(H, \Gamma) = \widetilde{H}/_{\tau_\partial}$, where τ_∂ is the restriction of τ to $\partial\widetilde{H} \setminus \widetilde{\Gamma}$. When H is orientable, $D(H, \Gamma)$ is just $H \setminus \Gamma$ mirrored along its boundary $(\partial H) \setminus \Gamma$. In general, $D(H, \Gamma)$ is an open orientable manifold with cusps based on tori. If $T \in \mathcal{T}_n$ we can construct $D(T) := D(N(T))$, and define $\mathcal{D}_n$ as the set of manifolds of this type. We also set $\mathcal{D} = \bigcup_{n=1}^{\infty} \mathcal{D}_n$.

The manifold $D(T)$ is constructed in [4] as the boundary of the 4-dimensional thickening of a certain 2-dimensional polyhedron $P_0(T)$ "dual" to T (this construction is summarized below in Section 3). The analogue of Corollary 0.3 for $\mathcal{D}$ is also proved in [4], and it says that every closed orientable 3-manifold is a Dehn filling of a manifold in $\mathcal{D}$, *i.e. the class* $\mathcal{D}$ *is universal for* 3*-manifolds under Dehn filling.* The following result shows that the elements of $\mathcal{D}$ correspond bijectively to those of $\mathcal{T}$, so they can be easily classified in purely combinatorial terms.

Theorem 0.4. *Every member of* $\mathcal{D}$ *is an orientable cusped hyperbolic manifold without boundary. The correspondence* $T \mapsto D(T)$ *defines a bijection between* $\mathcal{T}$ *and* $\mathcal{D}$*.*

We mention here that the set $\mathcal{D}$ consists of the manifolds which decompose into "blocks" homeomorphic to that shown in Fig. 1. This block is also used in [2] to construct a hyperbolic manifold from a path in the curve complex of a surface. In Section 3 we show that every manifold in $\mathcal{D}$ (with a very few exceptions) contains a closed incompressible genus-2 surface.

Isometries Since a Dehn filling of a hyperbolic manifold "typically" is hyperbolic, the class $\mathcal{D}$ provides a powerful method to construct closed hyperbolic manifolds. As an application of this method we establish in Section 4 the following result, the first statement of which was already known [13].

Proposition 0.5. *If* G *is a finite group then there exists a closed orientable hyperbolic* 3*-manifold* V *such that the isometry group of* V *is isomorphic to* G*. Moreover* V *can be chosen so that* $\mathrm{vol}(V) \leqslant c \cdot |G|^9$*, where* $c > 0$ *is a constant.*

Exceptional slopes In Section 3 we study some Dehn fillings of the manifolds in $\mathcal{D}$. To state our main results we recall that, according to Thurston's hyperbolic Dehn filling theorem, on each cusp of a finite-volume hyperbolic 3-manifold there is only a finite number of slopes filling along which one gets a non-hyperbolic 3-manifold. These slopes are called *exceptional*, and a considerable effort has been devoted to understanding them [11]. If T is a triangulation, the hyperbolic manifold $D(T)$ has a preferred horospherical cusp section, and each component of this section corresponds to an edge of T. Moreover the valence of the edge gives a lower bound for the length of the second shortest geodesic on the component. This fact and the Agol-Lackenby 6-theorem [1, 15] imply the following:

Proposition 0.6. *If every edge of T has valence at least 7 then there is at most one exceptional slope on each cusp of $D(T)$.*

Hyperbolicity from combinatorics Suppose a 3-manifold M has an *ideal* triangulation T such that each edge has valence at least 6. An easy argument shows that $\chi(T) \leqslant 0$, and that $\chi(T) = 0$ precisely when all valences are 6. Moreover, in the last case, the boundary of M is a disjoint union of tori and Klein bottles, and Thurston's hyperbolicity equations for cusped manifolds have a very simple solution, given by regular ideal tetrahedra. Analogously, if all edges have one and the same valence $v \geqslant 7$, the hyperbolicity equations for the geodesic boundary case described in [10] have a simple solution, given by regular truncated tetrahedra with dihedral angles $2\pi/v$. An argument based on the Agol-Lackenby machinery [1, 15] allows us to generalize these facts as follows (see Section 3):

Proposition 0.7. *If M has an ideal triangulation T whose edges have valence at least 6, then M is hyperbolic, and the edges of T are homotopically non-trivial relative to ∂M.*

If all valences are at least 6, every edge of T is then homotopic to a geodesic. This shows that the tetrahedra of T themselves are homotopic to straight truncated ones, so it seems quite natural to ask whether such tetrahedra are in fact *isotopic* to straight ones, *i.e.* whether T gives a geometric decomposition of $M(T)$. This is true when the valences of the edges are all equal to each other, as discussed above. And we have verified experimentally the same fact in [9] for all triangulations involving at most 4 tetrahedra. We therefore propose the following conjecture:

Conjecture 0.8. *If M has an ideal triangulation T whose edges have valence at least 6, then T is realized by hyperbolic partially truncated tetrahedra.*

From normal to essential surfaces Since understanding incompressible surfaces in arbitrary 3-manifolds is in general rather difficult, whereas normal surfaces can be studied algebraically (as solutions of Diophantine equations), the reader may find the next result of some interest. We define a surface in a relative handlebody (H, Γ) to be *essential* when its boundary consists of loops parallel to loops in Γ, it is incompressible, ∂-incompressible near Γ, and not ∂-parallel. Our definition is different from that given in [12]: see Section 1 for details.

Proposition 0.9. *Let $T \in \mathcal{T}$ be a triangulation of a 3-manifold. Then the essential surfaces in $N(T)$ correspond bijectively to the non-trivial normal surfaces in T.*

It was proved in [8] that an ideal triangulation with one edge does not contain any non-trivial normal surface. Triangulations with one edge correspond by Theorem 0.1 to relative handlebodies (H, Γ) with Γ consisting of a single hyperbolic loop of minimal complexity. The previous proposition then implies:

Corollary 0.10. *Let H be a handlebody and Γ be a two-sided hyperbolic loop of minimal complexity $c(\Gamma)$ on ∂H. Then there are no essential surfaces in (H, Γ).*

Acknowledgments We thank Simon King for communicating to us a proof, based on traditional cut-and-paste techniques, of Corollary 0.2 for the special case of genuine triangulations (without multiple and self-adjacencies). F. C. thanks Dylan Thurston for his continuous encouragement and illuminating observations. B. M. thanks the Mathematics Department of the University of Austin for hospitality during May 2003. C. P. thanks the Université Paris 7 and the Institut de Mathématiques de Jussieu for hospitality during May 2003.

1 Triangulations and hyperbolic relative handlebodies

In this section we give a precise definition of what we call a triangulation, and we show that the set of triangulations corresponds to a class of hyperbolic structures on relative handlebodies. We also provide geometric and topological information on these handlebodies, including two intrinsic definitions of their class.

Triangulations We call *triangulation* a pair $(\{\Delta_i\}_{i=1}^n, \{g_j\}_{j=1}^{2n})$, where $n \in \mathbb{N}$ is positive, each Δ_i is a copy of the standard tetrahedron, and the g_j's give a complete system of simplicial pairings between the faces of the Δ_i's, such that the gluing of the Δ_i's along the g_j's is connected. Note that this gluing may actually not be a 3-manifold, because the link of a vertex could be any surface, and the link of the midpoint of an edge could be a projective plane. Even when the gluing is a manifold, the Δ_i's give a triangulation only in a loose sense, because multiple adjacencies and self-adjacencies are allowed. Triangulations are viewed up to combinatorial equivalence, and the set of equivalence classes is denoted by $\mathcal{T}$.

If Δ is the tetrahedron, we denote now by Δ^* the polyhedron obtained by truncating Δ at the vertices (*i.e.*, formally, by removing open stars of the vertices). Note that $\partial\Delta^*$ consists of four truncation triangles and four "lateral" hexagons. If $T = (\{\Delta_i\}_{i=1}^n, \{g_j\}_{j=1}^{2n})$ is a triangulation then the g_j's give gluing rules for the lateral hexagons of the Δ_i^*'s, and we denote by $M(T)$ the result, which may or not be a manifold. In case $M(T)$ is a manifold, T is called an *ideal triangulation* of $M(T)$. Note that in this case $\partial M(T)$ is triangulated (in a loose sense) by the truncation triangles of the Δ_i^*'s.

From triangulations to relative handlebodies Fix a triangulation $T = (\{\Delta_i\}_{i=1}^n, \{g_j\}_{j=1}^{2n})$. If we remove from each Δ_i an open neighbourhood of the whole 1-skeleton, we can still glue what is left along the restrictions of the g_j's. This corresponds to adding 1-handles to a disjoint union of 0-handles, so the result is a (possibly non-orientable) handlebody H. Of course H can also be obtained from $M(T)$ by removing an open neighbourhood of the image of the set of edges, as discussed in the Introduction. Each component of this neighbourhood is a cyclic gluing of wedges, which implies that $\partial H \setminus \partial M(T)$ is a union of strips (annuli and/or Möbius bands). We can then consider the system Γ of the cores of these strips, and get a relative handlebody $(H, \Gamma) =: N(T)$. Note that $M(T)$ is a manifold if and only if a regular neighbourhood of Γ in ∂H consists of annuli only (no Möbius strips).

Hyperbolic structure We have already defined in the Introduction what we mean by *hyperbolic structure* on a relative handlebody. We now explain how to construct one such structure on each $N(T)$.

Given a manifold with a triangulation $T = (\{\Delta_i\}_{i=1}^n, \{g_j\}_{j=1}^{2n})$, there is a general strategy to algorithmically construct a hyperbolic structure on $N(T)$ using T. This strategy, which dates back to Thurston [17] and was

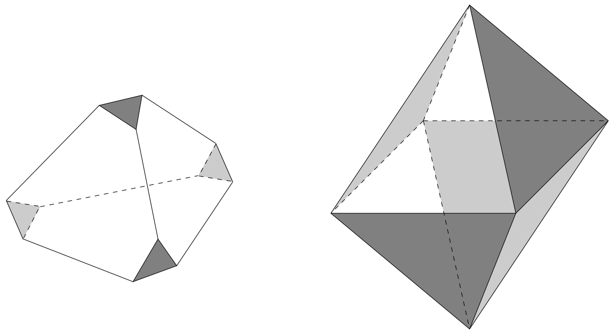

Figure 2: A truncated tetrahedron with dihedral angles zero is a regular ideal octahedron.

explained in detail in [10] for the geodesic boundary case, amounts to choosing the dihedral angles of each Δ_i along its edges, which gives Δ_i the shape of a hyperbolic polyhedron in $\mathbb{H}^3$ with some points at infinity, and then requiring that the hyperbolic polyhedra glue up coherently to give a complete structure.

We do not need to reproduce the details of this strategy here. We only note that, when trying to construct a hyperbolic structure on a relative handlebody $N(T)$, all the dihedral angles in T are actually forced to be 0. So in this case there is no choice to make: we only need to analyze what is a tetrahedron with 0 dihedral angles, and to show that gluing such objects we get a coherent and complete structure.

To accomplish the first task, start with a truncated tetrahedron as in Fig. 2-left, with black truncation triangles and white lateral hexagons. Now recall that declaring the dihedral angle along an edge to be 0 geometrically means that the edge disappears into an ideal vertex. The geometric version of the tetrahedron is therefore an ideal octahedron, with a checkerboard coloring of the faces, as in Fig. 2-right. Each dihedral angle of the octahedron is $\pi/2$, because it is the angle between a (now ideal) truncation triangle and a lateral "hexagon," now degenerated into an ideal triangle. It follows that the octahedron is regular. We denote now by O the regular ideal octahedron with checkerboard coloring of the faces as in Fig. 2.

To understand the gluing geometrically, we note that only the white triangles get glued (in pairs), while the black triangles give the geodesic boundary. Coherence of the structure along the gluings is not an issue, because any two ideal triangles are isometric to each other. Moreover, in the language of [10], there are no internal edges, so the gluing automatically gives a (possibly incomplete) hyperbolic structure on $N(T)$. However, completeness is automatic when there are no ends based on tori [10], so indeed $N(T)$ is hyperbolic. This already implies Corollary 0.3.

To state our next (easy) result we must recall that if V is hyperbolic with geodesic boundary and we consider a cusp of V based on a strip, then the strip has a well-defined Euclidean structure up to rescaling. If $(H,\Gamma) = N(T)$ and $\gamma \in \Gamma$ we can then speak of the Euclidean structure of the strip $U(\gamma)$ on which the cusp corresponding to γ is based. We also recall that γ comes from an edge of T.

Lemma 1.1. *If γ comes from an edge of valence q then the Euclidean structure of $U(\gamma)$ is obtained from a rectangle $[0,q]\times[0,1]$ by gluing $\{0\}\times[0,1]$ to $\{q\}\times[0,1]$.*

Proof. The cross-section of the octahedron O at a vertex is a square, with two opposite black edges and two white edges. As we construct $U(\gamma)$ we glue together q such squares along white edges. $\square$

Canonical decomposition In the projective model of $\mathbb{H}^3$ consider a polyhedron P with some ideal and some ultra-ideal vertices, and assume that all the edges of P meet the closure of $\mathbb{H}^3$. Dual to the ultra-ideal vertices of P there are hyperplanes of $\mathbb{H}^3$, truncating P along which we get the so-called "partially truncated polyhedron" associated to P. Kojima [14] has shown that a finite-volume hyperbolic 3-manifold V with non-empty geodesic boundary has a canonical decomposition $\mathcal{K}(V)$ into partially truncated polyhedra. Combinatorially, $\mathcal{K}(V)$ corresponds to a gluing $(\{P_i\},\{g_j\})$ of genuine polyhedra —an obvious extension of our notion of triangulation, where we allow arbitrary polyhedra instead of tetrahedra— and we use the same symbol $\mathcal{K}(V)$ to denote both the geometric and the combinatorial version of the decomposition. When the cusps of V are based on strips only, $\mathcal{K}(V)$ is actually dual to the cut-locus of ∂V. We can now establish the following:

Proposition 1.2. *If $T\in\mathcal{T}$ then $\mathcal{K}(N(T))$ consists of the regular ideal octahedra employed to construct the hyperbolic structure. In particular, $\mathcal{K}(N(T))$ equals T combinatorially.*

Proof. Since $\mathcal{K}(N(T))$ is dual to the cut-locus of $\partial N(T)$, it suffices to prove that the latter is obtained by gluing together the cut-loci of the black faces within the individual octahedra O employed to construct $N(T)$. The cut locus of the black faces of O is shown in Fig. 3-left combinatorially, and in Fig. 3-centre geometrically.

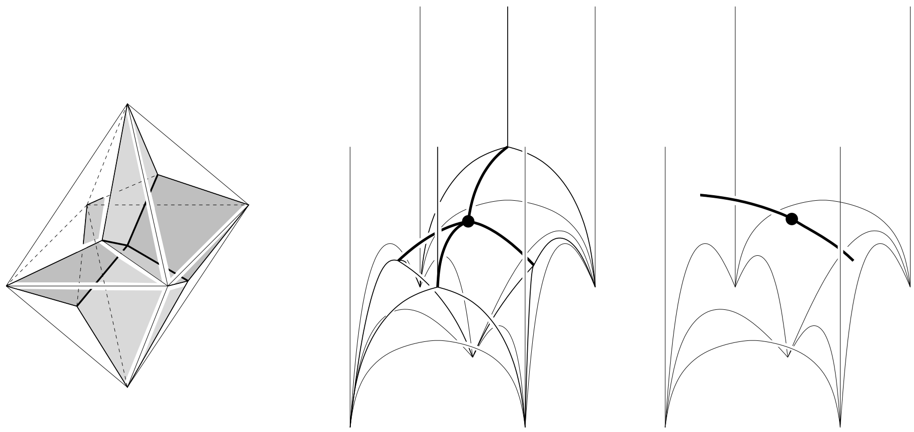

Figure 3: Cut locus of the black faces of O, and shortest path from a black to a white face.

To prove our result we must show that if $p \in N(T)$ and q is a point of $\partial N(T)$ closest to p then p and q belong to the same copy of O. If this is not the case, the shortest geodesic from q to p starts with an arc which leaves a black face of a copy of O orthogonally at q and exits from a white face. The shortest such arc joins the centre of a black face to the centre of the opposite white face, as shown in Fig. 3-right. Comparing with Fig. 3-centre we see that the arc is at least twice as long as the distance from p to the black faces of the copy of O in which p lies. This of course gives a contradiction. $\square$

An alternative proof of the previous proposition could also be given by computing the *tilts* of the faces of T, as explained in [10]. The proposition readily implies that the map $T \mapsto N(T)$ is a bijection between $\mathcal{T}$ and a set $\mathcal{H}$ of hyperbolic relative handlebodies (the main assertion of Theorem 0.1). Moreover $\mathcal{H}$ consists precisely of the hyperbolic manifolds which decompose into copies of the octahedron O, with black faces on the boundary and white faces in the interior.

The same result also allows us to prove that a triangulation of a 3-manifold Y is determined by its 1-skeleton as a subset of Y (Proposition 0.2). To state this precisely, denote by $\widehat{Y}$ the space obtained from Y by collapsing each component of ∂Y to a point. Then note that if we have a triangulation of Y with ideal vertices at the components of ∂Y and possibly some non-ideal vertices, then the 1-skeleton of the triangulation is a graph contained in $\widehat{Y}$, well-defined up to isotopy.

Proposition 1.3. *Let the triangulations T_0 and T_1 have the same 1-skeleton $J \subset \widehat{Y}$. Then T_0 and T_1 are isotopic relatively to J.*

Proof. For $i = 0, 1$ let $N(T_i) = (H_i, \Gamma_i)$, and observe that $\widehat{Y} \setminus J \cong H_i \setminus \Gamma_i$, so the identity map of $\widehat{Y} \setminus J$ induces a well-defined homeomorphism $h : (H_0, \Gamma_0) \to (H_1, \Gamma_1)$. Moreover, the geometric realization of the tetrahedra of T_i defines on $H_i \setminus \Gamma_i$ a hyperbolic metric. Mostow's rigidity theorem then gives a homotopy $\left(h_t : H_0 \to H_1\right)_{t\in[0,1]}$ such that $h_0 = h$, h_1 induces an isometry $H_0 \setminus \Gamma_0 \to H_1 \setminus \Gamma_1$, and for all $t \in [0, 1]$ the following happens:

$$h_t(\partial H_0) = \partial H_1, \quad h_t(\Gamma_0) = \Gamma_1, \quad h_t^{-1}(\Gamma_1) = \Gamma_0.$$

(These facts are not obvious *a priori.* They follow from the construction of the homotopy as a convex combination of two maps defined on the universal cover). Since T_i is the canonical Kojima decomposition of $H_i \setminus \Gamma_i$, we have $h_1(T_0) = T_1$. Suppose for a moment that each h_t is a homeomorphism and $h_t(\Gamma_0) = \Gamma_1$. Then we can use the natural identification $\widehat{Y} \setminus J \cong H_i \setminus \Gamma_i$ and extend h_t to a homeomorphism of $\widehat{Y}$ such that $h_t(J) = J$, thus getting the desired isotopy between T_0 and T_1. The rest of the proof is devoted to proving that $\left(h_t\right)_{t\in[0,1]}$ can indeed be replaced by an isotopy mapping Γ_0 to Γ_1 for all t.

A result of Waldhausen [19, Theorem 7.1] implies that there exists an isotopy $\left(h'_t : H_0 \to H_1\right)_{t\in[0,1]}$ such that $h'_0 = h_0 = h$ and $h'_1 = h_1$. Note that the trace f'_t of h'_t on ∂H_0 need not map Γ_0 to Γ_1, so $\left(h'_t\right)_{t\in[0,1]}$ is not the desired isotopy yet. To modify it, we return to the initial $\left(h_t\right)_{t\in[0,1]}$ and for $j = 0, 1$ we define $f_j : \partial H_0 \to \partial H_1$ as the restriction of h_j. Now recall that surface homeomorphisms which are homotopic relatively to the boundary are in fact isotopic [6]. Using this fact one easily sees that f_0 and f_1 embed into an isotopy $\left(f_t : \partial H_0 \to \partial H_1\right)_{t\in[0,1]}$ such that $f_t(\Gamma_0) = \Gamma_1$ for all $t \in [0, 1]$.

We have constructed two paths $\left(f'_t\right)_{t\in[0,1]}$ and $\left(f_t\right)_{t\in[0,1]}$ of homeomorphisms $\partial H_0 \to \partial H_1$, with the same ends. Now recall that the space of isotopically trivial automorphisms of a closed hyperbolic surface is simply connected [5]. Therefore, if we denote by H'_i the handlebody obtained from H_i by attaching an external collar, we can attach to $h'_t : H_0 \to H_1$ a product map between the collars, getting a certain $h''_t : H'_0 \to H'_1$ whose trace on the boundary is f_t.

Since $f_t(\Gamma_0) = \Gamma_1$, up to identifying H'_i to H_i, the isotopy $\left(h''_t\right)_{t\in[0,1]}$ is now the desired one. To be completely formal and make sure that $h''_j = h_j$ for $j = 0, 1$ one should at the beginning slightly perturb h_0 and choose

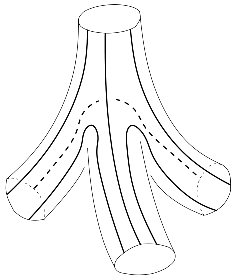

Figure 4: A tetrapod.

internal collars of ∂H_i on which h_0 and h_1 behave as products, and then identify H_i' to H_i just by rescaling the two-sided collar of ∂H_i in H_i' to its one-sided collar in H_i. $\square$

Hyperbolic volume Of course each element of $\mathcal{H}_n$ (a relative handlebody of the form $N(T)$ with $T \in \mathcal{T}_n$) has volume $n \cdot v_O$, where v_O is the volume of the hyperbolic regular ideal octahedron. Moreover $v_O = 8\Lambda(\pi/4) \approx 3.66386$, as first computed by Milnor in [17].

If $T \in \mathcal{T}_n$ then the numbers of 0- and 1-handles used to construct $N(T)$ are respectively n and $2n$, so $N(T)$ has genus $n+1$ and its boundary has Euler characteristic $-2n$. Miyamoto [16] has shown that $\mathcal{H}_n$ is precisely the set of hyperbolic manifolds Y having minimal volume among those with $\chi(\partial Y) = -2n$.

Corollary 1.4. *$\mathcal{H}_n$ consists of the hyperbolic relative handlebodies of minimal volume among those based on a handlebody of genus $n+1$.*

Complexity We now turn to the pant-meridinal complexity defined in the Introduction, establishing all the remaining assertions of Theorem 0.1.

If $T \in \mathcal{T}_n$ and $N(T) = (H, \Gamma)$, consider first the $2n$ meridinal discs of the 1-handles used to construct H. These discs cut H into a union of "tetrapods" as shown in Fig. 4. Note that Γ meets the boundary of each disc in three points. To get a decomposition as in the definition of $c(\Gamma)$ we need to cut each tetrapod along another disc. There are three non-isotopic but equivalent discs one could choose, and each meets Γ in 4 points. This decomposition gives the estimate $c(\Gamma) \leqslant 10n$.

We now claim that if (H, Γ) is hyperbolic and H has genus $n+1$ then $c(\Gamma) \geqslant 10n$. Assume the contrary and pick a family of discs which decomposes H into solid pairs of pants and meets Γ in less than $10n$ points. Then there must exist one solid pair of pants whose three meridinal discs meet Γ in $k < 10$ points. But k is even, so $k \leqslant 8 < 9 = 3 \times 3$, whence there is one disc meeting Γ in $h \leqslant 2$ points. Depending on whether h is 0, 1, or 2, this disc gives either a boundary-compressing disc in (H, Γ), or a disc compressing a toric (or Klein bottle) cusp in the genuine double of (H, Γ), or an essential annulus in the double of (H, Γ). Since (H, Γ) and its double are hyperbolic, we get a contradiction.

We are left to show that a hyperbolic relative handlebody (H, Γ) with H of genus $n+1$ and $c(\Gamma) = 10n$ arises as $N(T)$ for some triangulation T. Let us consider the decomposition of H into solid pairs of pants which realizes the minimum $c(\Gamma)$. By what already proved, for every pair of pants in the decomposition there are two discs meeting Γ in three points, and one disc meeting Γ in four points. It is now easy to show that there is only one such configuration such that the number of intersections between Γ and the discs cannot be reduced, whence the conclusion at once.

Essential surfaces We prove here Proposition 0.9. Let M be a manifold with an ideal triangulation T, and let Σ be a closed surface contained in the interior of M. Haken defined Σ to be *normal* if it is transversal to T and it meets each tetrahedron of T in a union of squares and triangles. The set of normal surfaces in T can be described combinatorially as the set of "acceptable" solutions of a system of linear Diophantine equations, with unknowns corresponding to the numbers of triangles and squares in each tetrahedron, and equations imposing consistency of the number of intersections with each edge. A solution is acceptable if each tetrahedron contains one type of square only. A normal surface is *trivial* when it consists of triangles only, which implies that it is parallel to the boundary.

Given a relative handlebody (H, Γ), we denote by $U(\Gamma)$ a regular neighbourhood of Γ in ∂H and we define $\mathcal{S}(H, \Gamma)$ as the set of properly embedded surfaces $\Sigma \subset H$ such that all the components of $\partial\Sigma$ are contained in $U(\Gamma)$ and essential in it. The elements of $\mathcal{S}(H, \Gamma)$ are viewed up to isotopy through surfaces of the same type. For the elements of $\mathcal{S}(H, \Gamma)$ we use the customary notion of *incompressibility.* We also say that $\Sigma \in \mathcal{S}(H, \Gamma)$ is *∂-incompressible near* Γ if every ∂-compressing disc D with $\partial D \cap \partial H \subset U(\Gamma)$ is trivial. We then define a connected $\Sigma \in \mathcal{S}(H, \Gamma)$ to be *essential in* (H, Γ) if it is incompressible, ∂-incompressible near Γ, and not parallel to a component of $\partial H \setminus \Gamma$.

In the next results we fix $T \in \mathcal{T}$ and set $(H, \Gamma) = N(T)$. We also assume that $M(T)$ is a manifold, so $U(\Gamma)$ consists of annuli (but H may be non-orientable). Under this assumption one sees that $\mathcal{S}(H,\Gamma)$ can be identified to the set of closed surfaces contained in the interior of $M(T)$ and transversal to $T^{(1)}$, up to isotopy through surfaces of the same type.

Lemma 1.5. *If two normal surfaces coincide as elements of $\mathcal{S}(H,\Gamma)$ then they coincide as normal surfaces.*

Proof. Isotopy within $\mathcal{S}(H,\Gamma)$ preserves the intersection numbers with the edges of T, which determine the surface. $\square$

Theorem 1.6. *Let $\Sigma \in \mathcal{S}(H,\Gamma)$ be connected. Then Σ is essential if and only if it is normal and non-trivial, up to isotopy.*

Proof. We prove that Σ is normal if and only if it is incompressible and ∂-incompressible near Γ. The conclusion then follows, since ∂-parallel surfaces correspond to trivial normal surfaces.

Assume Σ is normal. We first prove it is incompressible. Suppose $D \subset H$ is a disc compressing Σ. Note that D can be viewed as a subset of $M(T)$, and it does not touch $T^{(1)}$. We can assume that D is transversal to $T^{(2)}$, so $D \cap T^{(2)}$ is a union of proper arcs and circles in D. Using irreducibility of H one easily sees that the circles can be dismissed by an isotopy of D. Now consider a proper arc α which is outermost in D, and note that there is a portion D' of D such that $\partial D' = \alpha \cup \beta$ and $\beta \subset \Sigma$. Moreover D' lies within one tetrahedron Δ_i of T. By assumption $\Sigma \cap \Delta_i$ is a union of triangles and squares. A triangle or a square intersects each face of Δ_i once, so we can isotope D in such a way that β slides over Σ and α disappears from $D \cap T^{(2)}$. Iterating this argument we can assume $D \cap T^{(2)}$ is empty, so D lies within a tetrahedron of T, hence it is trivial: a contradiction.

We now prove Σ is ∂-incompressible near Γ. By contradiction, let $D \subset H$ be a disc such that $\partial D = \alpha \cup \beta$ with $\alpha = D \cap \partial H \subset U(\Gamma)$ and $\beta = D \cap \Sigma$, and note that the two components of $\partial\Sigma$ touched by D are distinct, for otherwise D would give a disc compressing the interior of Σ. Therefore, up to isotopy we can identify D to a subset of $M(T)$ in such a way that α is contained in an edge of T and a regular neighbourhood of α in D is contained in one tetrahedron Δ_i of T. We can then apply the above argument to eliminate all the components of $D \cap T^{(2)}$ except α. Therefore we can assume that $D \subset \Delta_i$, which is clearly impossible since no component of $\Sigma \cap \Delta_i$ meets the same edge twice.

Now let Σ be incompressible and ∂-incompressible near Γ. Since Σ is incompressible (and H is irreducible), Σ can be isotoped to a surface transverse to $T^{(1)}$ and intersecting each Δ_i in discs D_k^i with $D_k^i \cap T^{(1)} \neq \emptyset$. We now prove that each D^i_k is either a triangle or a square. If not, it meets an edge e of T in at least two points. Taking two consecutive points on e, the subarc α of e which joins them, and an arc β which joins them within D^i_k, we can find a ∂-compressing disc bounded by $\alpha \cup \beta$. This disc must be trivial, so the two components of $\partial\Sigma$ involved are actually the same, while the ends of α determine distinct components. □

Corollary 1.7. *An essential surface in $\mathcal{S}(H,\Gamma)$ is determined (up to isotopy) by its trace on ∂H.*

2 Mirrored relative handlebodies

In this section we study the class $\mathcal{D}$ of hyperbolic 3-manifolds obtained by the mirroring construction applied to the elements of $\mathcal{H} = N(\mathcal{T})$, as described in the Introduction. The class $\mathcal{D}$ was first defined in [4].

Spaces associated to triangulations There is a common idea underlying the construction of the sets $\mathcal{H}$ and $\mathcal{D}$ of manifolds in one-to-one correspondence with the set $\mathcal{T}$ of triangulations. This idea is to choose a topological or geometric block B with four distinguished subsets called *faces*, and to associate to a triangulation T a space $\Theta(T,B)$ obtained by gluing along faces some copies of B according to the combinatorics of T. For several choices of B we can then prove that the topology or geometry of $\Theta(T,B)$ determines T uniquely.

It turns out that, to realize the idea just sketched, an additional structure is necessary besides the block B and its four faces. Namely, we suppose that B comes equipped with a set of 24 automorphisms which act as $\mathfrak{S}_4$ on the faces. (This set may or not be a group, as we will see below). The basic remark for the construction of $\Theta(T,B)$ is then that, as in SnapPea's notation [20], a gluing in T can be viewed as a restriction of an automorphism of the tetrahedron Δ, whence of an element σ of $\mathfrak{S}_4$, so the given automorphism of B acting as σ allows to correspondingly glue the faces of B.

To be precise, we assume the faces of B are numbered as $F_0,\ldots,F_3$ and that the structure of B includes a map $f : \mathfrak{S}_4 \to \mathrm{Aut}(B)$ such that $f(\sigma)(F_i) = F_{\sigma(i)}$ and

$$f(p\cdot\sigma\cdot q) = f(p^{-1})^{-1}\circ f(\sigma)\circ f(q). \tag{1}$$

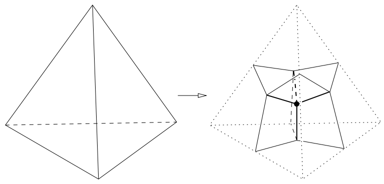

Figure 5: From a triangulation to the dual special polyhedron.

Let us then explain how to construct $\Theta(T,B)$ for a triangulation $T = \big(\{\Delta_i\}_{i=1}^n, \{g_j\}_{j=1}^{2n}\big)$. Number the faces of Δ as $\theta_0,\ldots,\theta_3$ and fix correspondences $\varphi_i : \Delta \to \Delta_i$. Let g_j glue $\varphi_{i_0(j)}(\theta_{k_0(j)})$ to $\varphi_{i_1(j)}(\theta_{k_1(j)})$. Then g_j determines $\sigma_j \in \mathfrak{S}_4$ such that g_j is the restriction to $\varphi_{i_0(j)}(\theta_{k_0(j)})$ of $\varphi_{i_1(j)} \circ \sigma_j \circ \varphi_{i_0(j)}^{-1}$, where σ_j also denotes the automorphism of Δ mapping θ_k to $\theta_{\sigma_j(k)}$. Now take n copies $\{B_i\}$ of B and fix correspondences $\widetilde{\varphi}_i : B \to B_i$. For all j glue $\widetilde{\varphi}_{i_0(j)}(F_{k_0(j)})$ to $\widetilde{\varphi}_{i_1(j)}(F_{k_1(j)})$ along the restriction of $\widetilde{\varphi}_{i_1(j)} \circ f(\sigma_j) \circ \widetilde{\varphi}_{i_0(j)}^{-1}$ and call $\Theta(T,B)$ the result. Relation (1), which is automatic when f is a homomorphism, ensures that $\Theta(T,B)$ is well-defined regardless of the choice of the φ_i's, as the reader can easily check.

We describe now four examples of blocks B. The map $f : \mathfrak{S}_4 \to \mathrm{Aut}(B)$ is a homomorphism in the first three cases, but not in the last case.

- Take Δ itself as B, with f the identity of $\mathfrak{S}_4 = \mathrm{Aut}(\Delta)$. Then $\Theta(T,\Delta)$ is the support $|T|$ of T;

- Let B be the cone over the 1-skeleton of the tetrahedron dual to Δ, as in Fig. 5, with the Y-shaped graphs playing the rôle of faces and the map f defined by restriction from $\mathfrak{S}_4 = \mathrm{Aut}(\Delta)$. Then $\Theta(T,B)$ is a *special polyhedron* $P(T)$, called the *dual* of T;

- Take as B the regular ideal octahedron O as in Fig. 2-right, with white triangles as faces, and f again given by restriction from $\mathfrak{S}_4 = \mathrm{Aut}(\Delta)$ (recall that O is obtained by truncating Δ). Note that the image of f consists of the isometries of O which preserve the colour of triangles. The resulting metric space $\Theta(T,O)$ is then isometric to $N(T)$, as shown in Section 1;

- Take as B the block K defined by mirroring O in the black triangles. This block was already considered in [2], and it was described above in Fig. 1. The white triangles of O glue up to 4 thrice-punctured spheres, which we define to be the faces of K. By Proposition 1.2, the isometry group of K is

$$\{\text{isometries of } O \text{ preserving colourings}\} \times \mathbb{Z}/_{2\mathbb{Z}} \cong \mathfrak{S}_4 \times \mathbb{Z}/_{2\mathbb{Z}}$$

where the non-trivial element of $\mathbb{Z}/_{2\mathbb{Z}}$ exchanges the two mirror copies of O. Identifying $\mathbb{Z}/_{2\mathbb{Z}}$ to $\{\pm 1\}$ we then define $f : \mathfrak{S}_4 \to \mathrm{Aut}(K)$ by $f(\sigma) = (\sigma, -\mathrm{sgn}(\sigma))$. Since the image of f consists of the orientation-reversing isometries of K, the resulting space $\Theta(T, K)$ is orientable, and it is not hard to check that it is actually isometric to the manifold $D(T)$, the orientable "double" of $N(T)$ defined in the Introduction.

It is a striking fact that, although $|T|$ does not determine T, the topology of the other spaces $P(T)$, $N(T)$, and $D(T)$ does indeed determine T. This is easy for $P(T)$, it was shown in Section 1 for $N(T)$, and it is proved below for $D(T)$:

Theorem 2.1. *The map $D : \mathcal{T} \to \mathcal{D}$ is bijective.*

As a preliminary to the proof of this result, note that $D(T)$ is a cusped orientable finite-volume hyperbolic 3-manifold without boundary. Its volume is $2v_O$ times the number of tetrahedra of T, *i.e.* twice the volume of $N(T)$. The cusps of $D(T)$ and $N(T)$ correspond to the edges of T, and the geometry of a cusp of $D(T)$ is determined by the geometry of the corresponding cusp in $N(T)$, which is a Euclidean Möbius band or annulus. Now recall that in general, if V is an orientable hyperbolic 3-manifold with cusps, a *horospherical cusp neighbourhood* $\mathcal{O}$ of V is a union of disjoint open subsets of V, one in each cusp, which lift to open horoballs in $\mathbb{H}^3$. Note that $\partial\mathcal{O}$ is a family of immersed Euclidean tori which intersect each other and self-intersect at most tangentially. The following result is an easy consequence of Lemma 1.1.

Corollary 2.2. *Let $T \in \mathcal{T}$ and let $\mathcal{O}(T)$ be the horospherical cusp neighbourhood defined in Fig. 6. Then $\mathcal{O}(T)$ is maximal. Moreover, if a component E of $\partial\mathcal{O}(T)$ bounds a cusp which corresponds to an edge of T with valence q, then E is a Euclidean torus of area $2q$ and geometric shape as described in Fig. 7.*

For the proof of Theorem 2.1 and later in Section 4 we will need to refer to certain specific triangulations T_1, T_2, $T_3^{(k)}$, and $T_4^{(k)}$ for $k \geqslant 1$. The first two are defined in Fig. 8-left and centre, while $T_3^{(k)}$ (respectively, $T_4^{(k)}$) is obtained from the triangulation $T_5^{(k)}$ of Fig. 8-right by matching (a, b, c) to

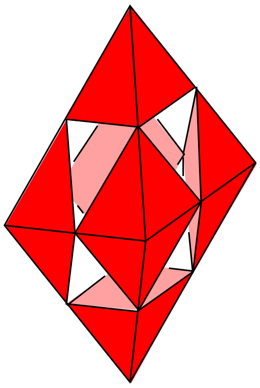

The decomposition of $D(T)$ into blocks K, and hence into octahedra, gives a horospherical cusp neighbourhood of $D(T)$, constructed by taking in each octahedron the horospherical neighbourhoods of the vertices shown here. Each such neighbourhood has volume equal to $1/2$ and is based on a unitary Euclidean square.

Figure 6: The horospherical cusp neighbourhood of $D(T)$ given by T.

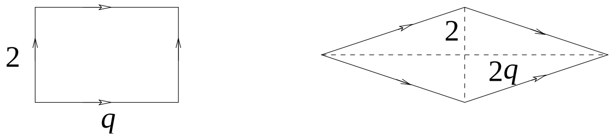

Figure 7: Shape of a boundary component of the horospherical cusp neighbourhood of $D(T)$. The rectangle occurs if the corresponding cusp of $N(T)$ is orientable, the rhombus if it is non-orientable.

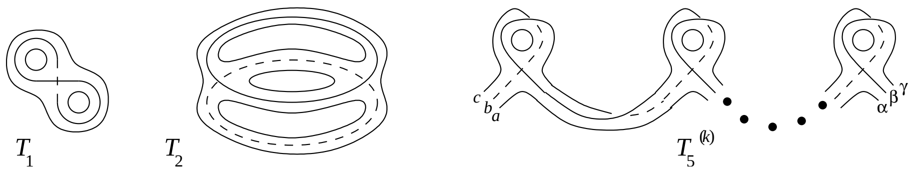

Figure 8: Particular triangulations. The picture describes the dual special polyhedron by means of the attaching circles of the regions to the singular set.

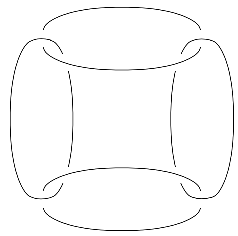

Figure 9: The minimally twisted chain link with 4 components.

(α, β, γ) (respectively, (α, γ, β)). Both $T_3^{(k)}$ and $T_4^{(k)}$ contain k tetrahedra. As an easy application of Corollary 2.2 one can now show the following:

Corollary 2.3. *For $j = 3, 4$ and $k \geqslant 1$ the manifolds $D(T_j^{(k)})$ are pairwise non-homeomorphic.*

We also note that $D(T_1)$ is a rather remarkable manifold. For instance, one can see that it is the complement in S^3 of the link shown in Fig. 9, and that its isometry group, which has 64 elements, acts as the dihedral group D_4 on the cusps.

The proof of Theorem 2.1 is an immediate consequence of the following:

Proposition 2.4. *Let V be an orientable cusped hyperbolic 3-manifold.*

1. *Every decomposition of V into blocks K arises from a unique triangulation T such that $V = D(T)$;*

2. *If $V \notin \{D(T_1), D(T_2), D(T_3^{(k)}), D(T_4^{(k)}) : k \geqslant 1\}$ then the decomposition of V into blocks K is unique;*

3. *If $V \in \{D(T_1), D(T_2), D(T_3^{(k)}), D(T_4^{(k)}) : k \geqslant 1\}$ then V has distinct decompositions into blocks K, but they arise from combinatorially equivalent triangulations.*

Proof of 2.4. The first assertion is easy: in a decomposition of V into blocks K all the gluings are orientation-reversing isometries, and the set of these isometries corresponds bijectively to $\mathfrak{S}_4 \cong \mathrm{Aut}(\Delta)$. Therefore the procedure described above to pass from T to $\Theta(T, K) = D(T)$ can be reversed, and we are done.

We establish assertions 2 and 3 at the same time, the proof of the latter being an easy by-product of the proof of the former. We give two different

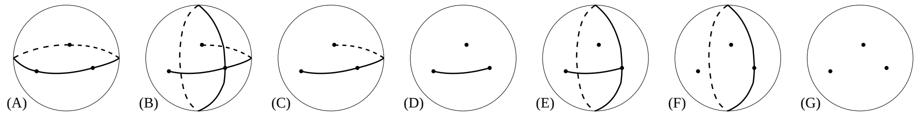

Figure 10: Disjoint systems of geodesics in the thrice-punctured sphere.

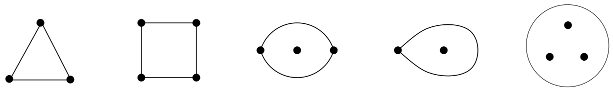

Figure 11: Surfaces obtained cutting S' along λ.

arguments, one analyzing the positions of geodesic thrice-punctured spheres in V, and one using the Epstein-Penner decomposition.

For the first proof, assume that V has two distinct decompositions ρ and ρ' into blocks K, arising from triangulations T and T'. We must show that T and T' are isomorphic to each other and to one of $T_1, T_2, T_3^{(k)}, T_4^{(k)}$, $k \geqslant 1$. Note first that T and T' must have the same number n of tetrahedra. For the sake of simplicity we first assume $n > 1$. Later we will sketch how to treat the easy case of triangulations with a single tetrahedron.

Remark now that the decompositions ρ and ρ' are made along $2n$ geodesic thrice-punctured spheres, and let $S_1, \ldots, S_{2n}$ be those giving ρ. Since $\rho' \neq \rho$ there is another such sphere S' distinct from the S_i's. Our first aim is to analyze how S' can intersect the S_i's. Before proceeding, note that the thrice-punctured sphere has a unique hyperbolic structure, and it contains precisely 6 complete simple geodesics.

Since S' and $\bigcup S_i$ are geodesic surfaces, they intersect transversely in a disjoint union λ of geodesic lines. Therefore λ as a subset of S' gives one of the configurations (A)-(G) shown in Fig. 10. Cutting S' along λ we then get one or two of the surfaces shown in Fig. 11, and each such surface must be contained in one of the blocks K arising from T. We now note that λ is a union of geodesics also in $\bigcup S_i$, and S_i is the double of one of the white faces of the octahedron O. It readily follows that the intersection between S' and each white triangle is either an edge or a height (the bisecant of a corner) of the triangle. Combining this fact with the remark that $S' \setminus \lambda$ contains at most two components as in Fig. 11, we see that the intersection between S' and each of the blocks K is empty or as in Fig. 12.

The components of $S' \cap K$ which appear in Fig. 12 are either ideal squares,

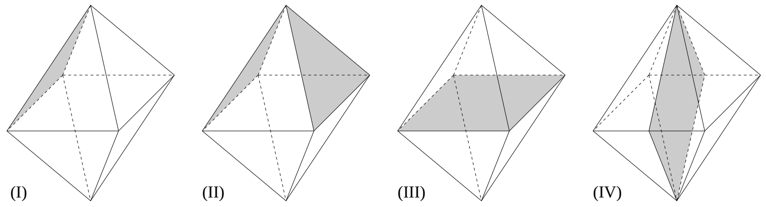

Figure 12: The intersection between S' and one of the blocks K. In all four cases we are only showing the octahedron O of which K is the double. In cases (I)-(III) the intersection $S' \cap K$ is contained in O, while in case (IV) it contains another similar quadrilateral in the mirror copy of O.

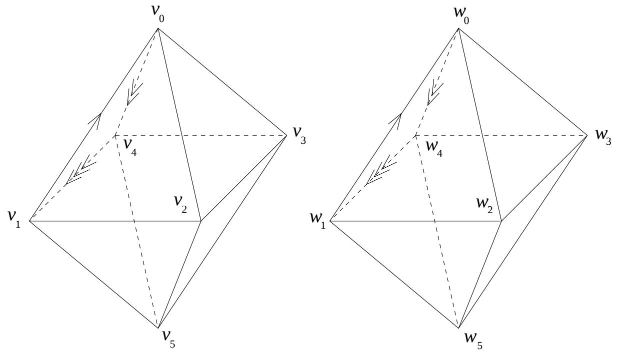

Figure 13: Configuration (A$'$).

or ideal triangles, or punctured ideal bigons. It follows that, among the configurations of λ on S' shown in Fig. 10, only cases (A)-(D) are possible. Moreover case (A) can be realized either by two triangles of type (I) in different blocks K, or by a pair of triangles of type (II) in one block. We call these configurations (A$'$) and (A$''$) respectively. In a similar way we have cases (B$'$) and (B$''$), while case (C) can only arise from one square of type (III), and case (D) only from a pair of quadrilaterals of type (IV).

We now prove that the existence of a thrice-punctured sphere S' of one of the types (A$'$), (A$''$), (B$'$), (B$''$), (C), or (D) forces a portion of T (or the whole of it) to have some definite shape. For instance, in case (A$'$) there must be two blocks K with gluings between their boundary thrice-punctured spheres that induce the edge-identifications of Fig. 13. Since gluings are orientation-reversing, the edge marked with a single arrow forces the gluing $(v_0, v_1, v_2) \to (w_0', w_1', w_2')$, where (v_0, v_1, v_2) is a white triangle of the

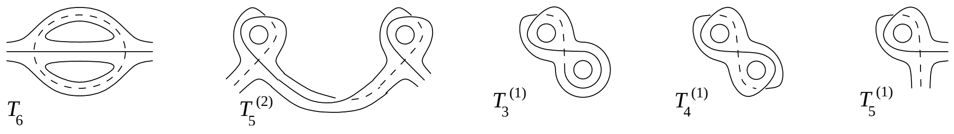

Figure 14: Particular triangulations or fragments of.

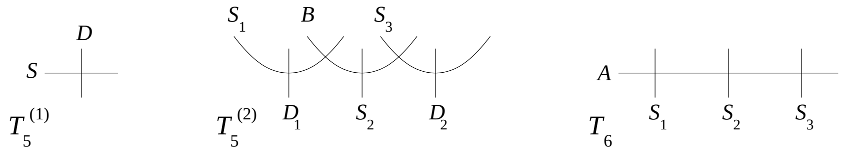

Figure 15: Patterns of intersections of surfaces in $\mathcal{S}$.

octahedron on the left, and (w_0', w_1', w_2') is the mirror image of (w_0, w_1, w_2) in the companion of the octahedron on the right. Similarly we must have gluings $(v_0, v_3, v_4) \to (w_0', w_3', w_4')$ and $(v_5, v_1, v_4) \to (w_5', w_1', w_4')$. It easily follows that T contains a subtriangulation T_6 as in Fig. 14. A long but straight-forward analysis shows that case (A$''$) leads to T_1 (defined above in Fig. 8), case (B$'$) to $T_5^{(2)}$, case (B$''$) to $T_3^{(1)}$ or $T_4^{(1)}$, case (C) to T_1 again, and case (D) to $T_5^{(1)}$. Cases T_1, $T_3^{(1)}$, and $T_4^{(1)}$ are actually forbidden by the assumption $n > 1$, so we discard them for the time being.

Having analyzed completely how a single geodesic thrice-punctured sphere S' can appear with respect to the S_i's, we now consider the collection $\mathcal{S}$ of all such spheres (including the S_i's). We begin by noting that the portions of triangulation found so far contain several elements of $\mathcal{S}$, namely T_1 contains two surfaces of type (A$''$)'s, four of type (C), and two S_i's, $T_3^{(1)}$ contains one (B$''$), one (D) and two S_i's, $T_4^{(1)}$ contains two (B$''$)'s, two (D)'s and two S_i's, $T_5^{(1)}$ contains one (D) and one S_i, $T_5^{(2)}$ contains one (B$'$), two (D)'s and three S_i's, and T_6 contains one (A$'$) and three S_i's. This implies first of all that $\mathcal{S}$ is finite. Moreover we can study how the various elements of $\mathcal{S}$ contained in the fragment intersect each other, which always occurs along one or more geodesics. The resulting configurations are shown by the 'incidence graphs' of Fig. 15 (where we omit T_1, $T_3^{(1)}$, and $T_4^{(1)}$, which are forbidden at this stage).

Let us consider now a connected component of $\bigcup \mathcal{S}$. By what shown already this component must be contained in a portion of T obtained by assem-

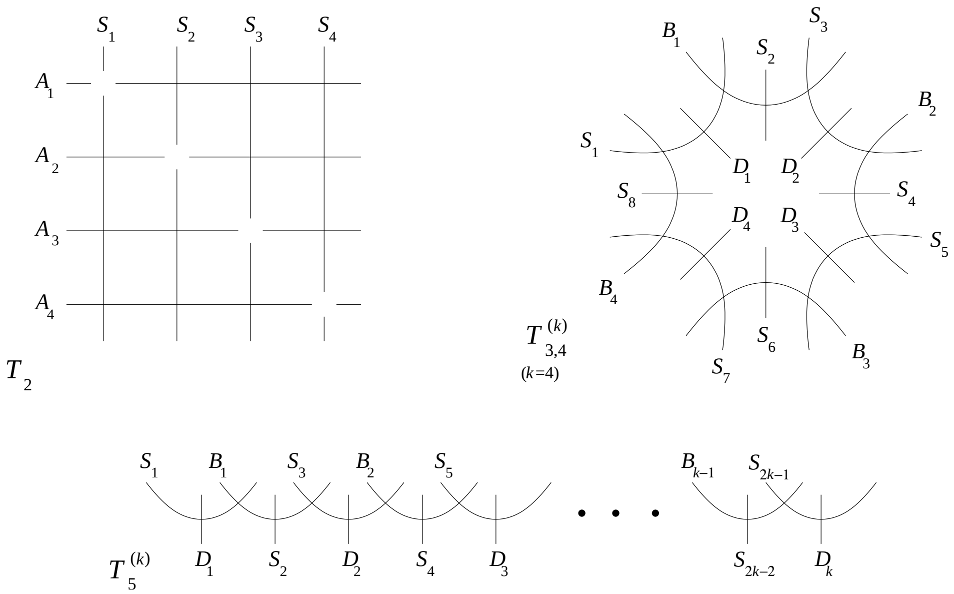

Figure 16: Patterns of intersections of elements of $\mathcal{S}$ in $T_2, T_3^{(k)}, T_4^{(k)}, T_5^{(k)}$.

bling fragments as in Fig. 14. This portion of T is then one of $T_2, T_3^{(k)}, T_4^{(k)}$ with $k \geqslant 2$, or $T_5^{(k)}$ with $k \geqslant 1$, or T_6. The incidence graphs of the elements of $\mathcal{S}$ contained in these triangulations are shown in Fig. 16 (except T_6, already shown above).

Let us note now that the incidence graphs which appear in Figg. 15 and 16 and can occur in V are pairwise distinct (remark that $T_3^{(k)}$ and $T_4^{(k)}$ have the same incidence graph, but V cannot be $D(T_3^{(k)})$ and $D(T_4^{(k)})$ at the same time, by Corollary 2.3). Recall that we are assuming that V has another decomposition ρ' into blocks K. The above discussion on how the set of all thrice-punctured spheres intersects the spheres giving ρ applies also to the spheres giving ρ'. We deduce that in the intersection graph of some component of $\mathcal{S}$ there should exist an automorphism not leaving the S_i's invariant. This can only happen in cases T_2, $T_3^{(k)}$, $T_4^{(k)}$, and $T_5^{(1)}$. However in $T_5^{(1)}$ the rôle of S_i can be identified intrinsically, because within it $S_i \cap D$ joins a puncture of S_i to itself, while it joins two distinct punctures of D. This eventually establishes assertion 2 for manifolds not decomposing in a single block K. To prove assertion 3 we only need to note that, as one sees from the incidence graphs, T_2, $T_3^{(k)}$, and $T_4^{(k)}$ contain precisely two families of punctured spheres giving a decomposition into blocks K, and these families arise from combinatorially equivalent triangulations.

We are left to deal with the case where $V = D(T)$ and T has one tetrahedron. As a by-product of the above argument we see that the decomposition into blocks K could be non-unique at most if T is T_1, $T_3^{(1)}$, or $T_4^{(1)}$. Now it is clear that $D(T_1)$, $D(T_3^{(1)})$, and $D(T_4^{(1)})$ are pairwise distinct, because they contain respectively 8, 6, and 4 thrice-punctured spheres. So at least the conclusion that $D(T)$ determines T is clear. A direct argument, again based on the analysis of the intersection graphs of the family $\mathcal{S}$, proves that in all these cases multiple decompositions indeed exist.

As mentioned in the Introduction, we have an alternative and more geometric proof of the proposition, which we believe is worth explaining. This proof is based on the polyhedral decomposition of $V = D(T)$ due to Epstein and Penner [7], which we now briefly recall. Let $\mathcal{O}$ be a horospherical cusp neighbourhood for V. If $\mathrm{Cut}(V, \mathcal{O})$ is the cut-locus of $V \setminus \mathcal{O}$ relative to its boundary, the Epstein-Penner decomposition of V relative to $\mathcal{O}$ is given by the ideal polyhedral decomposition of V dual to $\mathrm{Cut}(V, \mathcal{O})$. Let $\mathcal{O}(T)$ be as defined in Fig. 6. Let $\delta(T)$ denote the geometric decomposition of V into ideal octahedra induced by T. It is easy to see that $\delta(T)$ is the Epstein-Penner decomposition of V relative to $\mathcal{O}(T)$.

Let $\{e_i\}$ be the edges of T, and $\{q_i\}$ be the corresponding valences. Let C_i be the toric cusp of V corresponding to e_i. The component of $\mathcal{O}(T)$ corresponding to C_i has volume q_i. Observe that the total volume of $\mathcal{O}(T)$ is equal to $3 \cdot \#\delta(T) = 3 \cdot \mathrm{Vol}(V)/v_O$, so it depends solely on V.

Now let us suppose $V = D(T) = D(T')$ with $T \neq T'$. Let e'_i be the edge of T' corresponding to C_i and let q'_i be the multiplicity of e'_i in T'. If $q_i = q'_i$ for all i then $\mathcal{O}(T) = \mathcal{O}(T')$, so $\delta(T) = \delta(T')$ since $\delta(T)$ and $\delta(T')$ both give the Epstein-Penner decomposition of V relative to $\mathcal{O}(T) = \mathcal{O}(T')$. Moreover, the octahedra of $\delta(T)$ glue up in pairs into blocks K in a unique way, except when T is the triangulation T_2: two decompositions arise in this case, both induced by the same T anyway.

So let us suppose $q_i \neq q'_i$ for some i. Corollary 2.2 implies that one of the following conditions holds:

- $\{q_i, q'_i\} = \{1, 4\}$ and the cusps of $N(T)$ and $N(T')$ along e_i and e'_i are both orientable;

- $q_i = 2$, the cusp of $N(T)$ along e_i is orientable, $q'_i = 1$, and the cusp of $N(T')$ along e'_i is non-orientable; or viceversa.

Let v be a vertex of an octahedron O in $\delta(T)$. We define $r(v)$ to be the

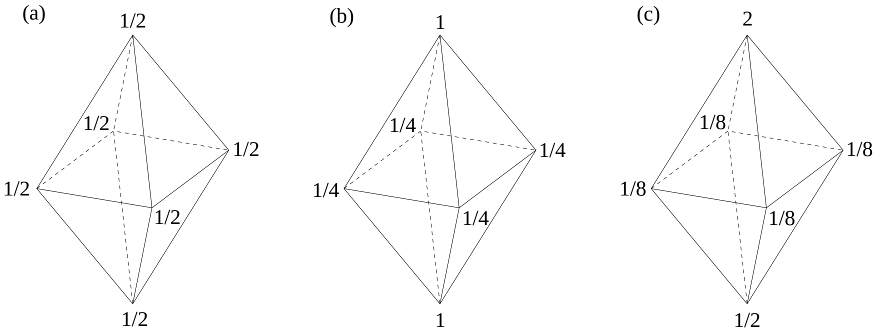

Figure 17: The allowed values for r on the vertices of an octahedron in T.

volume of the connected component of $\mathcal{O}(T') \cap O$ which is asymptotic to v. The previous observation implies that $r(v) \in \{1/8, 1/4, 1/2, 1, 2\}$. Note now that the components of $\mathcal{O}(T')$ lying within a certain octahedron $O \in \delta(T)$ do not overlap, and a direct computation shows that the volume of $\mathcal{O}(T') \cap O$ is then at most 3. Moreover the total volume of $\mathcal{O}(T')$ is thrice the number of octahedra of $\delta(T)$, so the volume of $\mathcal{O}(T') \cap O$ is precisely 3. This easily implies that every octahedron in $\delta(T)$ belongs to one of the three classes described in Fig. 17, according to the values that r takes on its vertices. If two vertices of octahedra of T are glued to each other then r takes on them the same value. Since V is connected we deduce that all the octahedra of T are of the same type.

If all the tetrahedra are of type (a) then $q_i' = q_i$ for all i, whence $T = T'$, as seen above. To analyze cases (b) and (c), note first that a vertex v with $r(v) = 2$ gives an edge of valence 1 in T, *i.e.* only one octahedron of $\delta(T)$ is incident to v, and the corresponding cusp of $N(T)$ is orientable. Similarly if $r(v) = 1/8$ then v gives an "orientable" edge of valence 4, if $r(v) = 1$ then v gives a "non-orientable" edge of valence 1, if $r(v) = 1/4$ then v gives an "orientable" edge of valence 2. Using these facts it is easy to see that if there is in $\delta(T)$ an octahedron of type (b) then there is only one, T is T_1, and T' is also T_1. Suppose then that all octahedra are of type (c). If there are k of them then T and T' are either $T_3^{(k)}$ or $T_4^{(k)}$, and Corollary 2.3 implies that $T = T'$ also in this case. $\square$

Proposition 2.4 enables us to compute the isometry group of $D(T)$ in terms of the group $\mathrm{Aut}(T)$ of combinatorial automorphisms of T.

Corollary 2.5. *Let T be a triangulation distinct from T_1, T_2, $T_3^{(k)}$, and $T_4^{(k)}$.*

Then

$$\mathrm{Isom}(D(T)) \cong \mathrm{Isom}^+(D(T)) \times \mathbb{Z}/_{2\mathbb{Z}},$$
$$\mathrm{Isom}^+(D(T)) \cong \mathrm{Isom}(N(T)) \cong \mathrm{Aut}(T).$$

Proof. The isomorphism $\mathrm{Isom}(N(T)) \cong \mathrm{Aut}(T)$ is an immediate consequence of Proposition 1.2. By construction every automorphism $\varphi \in \mathrm{Aut}(T)$ induces an orientation-preserving isometry $\overline{\varphi} \in \mathrm{Isom}(D(T))$. Moreover we can define an orientation-reversing isometry τ of $D(T)$ by reflecting any block K along the black faces of the octahedra the block is made of, and of course $\tau \circ \overline{\varphi} = \overline{\varphi} \circ \tau$ for any $\varphi \in \mathrm{Aut}(T)$. As shown in Proposition 2.4, if T is not $T_1, T_2, T_3^{(k)}, T_4^{(k)}$, then $D(T)$ admits a unique decomposition into blocks K, which is therefore preserved by any element in $\mathrm{Isom}(D(T))$. This implies that any isometry of $D(T)$ is induced, up to composition with τ, by an element in $\mathrm{Aut}(T)$, and our assertions easily follow. $\square$

3 Dehn filling

In this section we continue our study of the class $\mathcal{D}$ of complete hyperbolic 3-manifolds with cusps, introduced in [4] and shown in Section 2 to be in one-to-one correspondence with the set $\mathcal{T}$ of triangulations. It is proved in [4] that every element of $\mathcal{D}$ is a link complement in a connected sum of copies of $S^2 \times S^1$, and can be naturally seen as a subset of the boundary of a 4-manifold $X(T)$. In the same paper, this fact is used to prove that every 3-manifold bounded by a (possibly empty) set of tori is a Dehn filling of some $D(T)$.

New blocks We now describe three more spaces $\Theta(T, B)$ associated to a triangulation T by the choice of a block B with 4 faces and a preferred set of 24 automorphisms permuting the faces, as explained in Section 2. The third object is $X(T)$.

- Take as a block the set Q shown in Fig. 18-left, with the four Y's as faces. Since Q naturally embeds in the polyhedron of Fig. 5-right, we can as above define $f : \mathfrak{S}_4 \to \mathrm{Aut}(Q)$ by restriction from $\mathfrak{S}_4 = \mathrm{Aut}(\Delta)$. Then $\Theta(T, Q)$, that we denote by $P_0(T)$, is a *special polyhedron with boundary.* The boundary $\partial P_0(T)$ consists of loops that correspond to the edges of T. We note that $P_0(T)$ can be identified to a regular neighbourhood of the singular set of the special polyhedron $P(T)$ dual to T, defined in Section 2. So $P(T)$ is obtained from $P_0(T)$ by attaching a disc to each component of $\partial P_0(T)$;

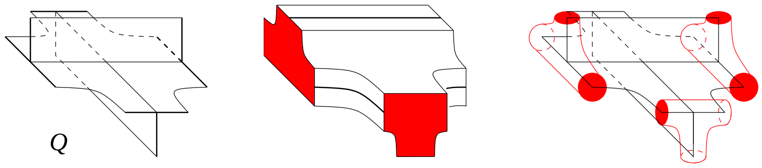

Figure 18: The three blocks giving respectively $P_0(T)$, $N(T)$, and $X(T)$.

- Let B be the pair $(\mathcal{N}_3(Q), \Gamma_Q)$, where $\mathcal{N}_3(Q)$ is the 3-dimensional thickening $\mathcal{N}_3(Q)$ of Q shown in Fig. 18-centre, and Γ_Q consists of 6 arcs in $\partial\mathcal{N}_3(Q)$, three of which are visible in the picture. We define the four shadowed polygons to be the faces, and we choose f again by restriction from $\mathfrak{S}_4 = \mathrm{Aut}(\Delta)$. Then $\Theta(T, B)$ is the relative handlebody $(H, \Gamma) = N(T)$. Note that $P_0(T) \subset H$ and $\Gamma = \partial P_0(T)$, and that H collapses onto $P_0(T)$;
- Take as a block the 4-dimensional thickening $\mathcal{N}_4(Q)$ of Q, defined by $\mathcal{N}_3(Q) \times [-1, 1]$. We choose as faces of $\mathcal{N}_4(Q)$ the products of the faces of $\mathcal{N}_3(Q)$ with $[-1, 1]$, each of which is a 3-ball collapsing to some Y-shaped face of Q, as suggested in Fig. 18-right. We define $f : \mathfrak{S}_4 \to \mathrm{Aut}(\mathcal{N}_4(Q))$ as we did to define $D(T)$, *i.e.* we set $f(\sigma)(x, t) = (\sigma(x), -\mathrm{sgn}(\sigma) \cdot t)$. The space $\Theta(T, \mathcal{N}_4(Q))$, that we denote by $X(T)$, is an *orientable* 4-dimensional manifold which contains $P_0(T)$ and collapses onto it.

Four-manifolds We begin with the following result, stated also in [4].

Proposition 3.1.
- *$\partial X(T)$ is a connected sum of copies of $S^2 \times S^1$;*
- *$\partial P_0(T)$ is a link in $\partial X(T)$;*
- *$D(T)$ is the complement of $\partial P_0(T)$ in $\partial X(T)$.*

Proof. Set $N(T) = (H, \Gamma)$. By construction $X(T)$ is the orientable $[-1, 1]$-bundle $H \tilde{\times} [-1, 1]$ over H. Since H is constructed using 0- and 1-handles only, the same holds for $X(T)$. This fact and the remark that $X(T)$ is orientable imply the first assertion. Now $\partial X(T)$ contains $\partial H \times \{0\}$ and ∂H contains $\Gamma = \partial P_0(T)$, which proves the second assertion. For the last one, we only need to note that in general the "orientable double" of a manifold V is defined as $\partial(V \tilde{\times} [-1, 1])$, and that $\partial X(T) \setminus \Gamma$ is homeomorphic to

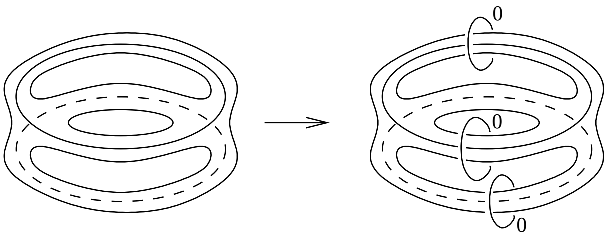

Figure 19: From a description of T via an immersion of $P_0(T) \subset S^3$ to a description of $D(T)$ via Dehn surgery.

$\partial\big((H \setminus \Gamma) \tilde{\times} [-1, 1]\big)$, because it is the same to remove Γ or a union of strips with cores in Γ. $\square$

Remark 3.2. The manifold $D(T)$ is not a link complement in S^3 in general. However, as noted in [4], we can always describe $D(T)$ as a link complement in S^3 with some 0-surgered unknots. To do this, we immerse $P_0(T)$ in S^3 as in the example of Fig. 19-left, we take the link $\partial P_0(T)$, and we encircle with 0-surgered unknots the triples of strands corresponding the complement of a maximal tree in the singular set of $P_0(T)$.

The polyhedron $P_0(T)$ is a *shadow* of the pair $(\partial X(T), \partial P_0(T))$. Shadows are defined in [18] for arbitrary links in 3-manifolds, and they can be used to show in particular that $\mathcal{D}$ is universal. We give a sketch of the proof of this fact, taken from [18, 4].

Theorem 3.3. *Every closed orientable 3-manifold is a Dehn filling of $D(T)$ for some $T \in \mathcal{T}$.*

Proof. It is a known fact that a closed, orientable 3-manifold V is the boundary of an orientable 4-manifold Y which admits a decomposition containing handles of indices 0, 1, and 2 only. The union of the cores of the handles of this handle decomposition is a 2-dimensional polyhedron P onto which Y collapses. Up to a slight perturbation we can assume that P is special, so it is a polyhedron $P = P(T)$ for some T. Recall that $P(T)$ is obtained from $P_0(T)$ by gluing discs to $\partial P_0(T)$. It is now sufficient to note that each time we glue a disc to a boundary component of $P_0(T)$ we are gluing a 2-handle to its unique orientable 4-dimensional thickening, hence, at the level of the boundary, we are performing a Dehn filling on a boundary component of the corresponding manifold $D(T)$. $\square$

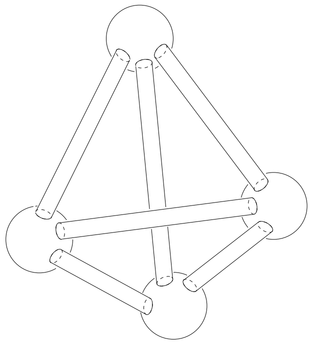

Figure 20: The block K is the complement in S^3 of the handlebody shown in the figure. The tubes give the cusps of K, and the thrice-punctured spheres (also viewed as pairs of pants) give the boundary of K, which is geodesic.

Remark 3.4. The above theorem was used in [4] as a starting point to define the *shadow-complexity* of a 3-manifold Y as the minimum number of vertices of a shadow for Y. The shadow-complexity of Y turns out to be strictly related to its geometry.

Meridians and longitudes Let T be an arbitrary triangulation. To provide some information on the Dehn fillings of $D(T)$, we first recall from Corollary 2.2 that $D(T)$ has a maximal horospherical cusp neighbourhood, bounded by a Euclidean torus of area $2q$ on each cusp. Here, q is the valence of the corresponding edge e of T, and the cusp has one of the two shapes shown in Fig. 7. The rectangle occurs if and only if the corresponding cusp in $N(T)$ is orientable, or equivalently if the regular neighbourhood of e in $M(T)$ is a manifold. Therefore, if T is an ideal triangulation of a manifold, only rectangles occur.

Recall that T induces a decomposition of $D(T)$ into blocks K. Each such block can be seen as the complement in S^3 of the handlebody shown in Fig. 20, and it contributes to the bases of the cusps with six Euclidean tubes of width 1 and boundary loops of length 2. We define any boundary loop of a tube to be the *meridian* of the corresponding cusp. In Fig. 7 the meridians are given by the vertical edges of the rectangle and by the vertical diagonal of the rhombus. If T is an ideal triangulation of a manifold, a preferred *longitude* on each cusp can also be defined as the unique loop having length q, *i.e.* as the loop given by the horizontal edges of the rectangle in Fig. 7-left.

Let T consist of n tetrahedra. It is clear from Fig. 20 that by performing a meridinal Dehn filling on each cusp we get a connected sum of $n+1$ copies of $S^2 \times S^1$, as stated in Proposition 3.1. Therefore the first Betti number of $D(T)$ is at least $n+1$. More precisely, we have the following:

Proposition 3.5. *Let $T \in \mathcal{T}$ consist of n tetrahedra. Then*

$$H_1(D(T); \mathbb{Z}) = \mathbb{Z}^{n+1} \oplus H^2(P(T); \mathbb{Z}).$$

Proof. Set $D = D(T)$ and $P = P(T)$. Denote by $S(P)$ the singular set of P, *i.e.* the 4-valent graph representing the gluings in T. Edges and vertices of $S(P)$ correspond to decomposing pairs of pants and blocks K of D. Let $p : D \to S(P)$ be the composition of the obvious projections $D \to N(T) \to P_0(T) \to S(P)$, and note that p sends a neighborhood of each decomposing pair of pants to the corresponding edge and each block K to the corresponding vertex. This gives a surjection $p_* : H_1(D) \twoheadrightarrow H_1(S(P)) = \mathbb{Z}^{n+1}$. Therefore $H_1(D) = \mathbb{Z}^{n+1} \oplus \mathrm{Ker}\,(p_*)$.

Every loop in $\mathrm{Ker}\,(p_*)$ is homologous to a sum of loops each contained in a block K. The cores of the six tubes of Fig. 20 generate $H_1(K)$, so the meridians μ_i, with arbitrary orientation, generate $\mathrm{Ker}\,(p_*)$. Each pair of pants separating two blocks K gives a relation of type $\pm\mu_{j_1} \pm \mu_{j_2} \pm \mu_{j_3} = 0$, with appropriate signs. We now claim that every other relation is superfluous.

A relation between the μ_i's comes from an embedded orientable surface bounded by some copies of the μ_i's. This surface intersects each pair of pants in a union of loops. The trivial loops can be dismissed by elementary cut-and-paste operations, so we can assume each loop is parallel to some μ_i. Therefore the relation can be expressed as a sum of relations coming from a set of surfaces each contained in a block K. Now consider K as the complement in S^3 of four balls $B_1, \ldots, B_4$ and six solid tubes, as in Fig. 20. Adding discs to a surface contained in K we get a closed surface in $S^3 \setminus (B_1 \cup \ldots \cup B_4)$. Any such surface compresses in $S^3 \setminus (B_1 \cup \ldots \cup B_4)$ to a union of spheres parallel to the ∂B_i's, so we can express the relation coming from a surface contained in K as a sum of relations induced by the pairs of pants in ∂K. This proves our claim.

We have shown that $\mathrm{Ker}\,(p_*)$ has one generator μ_i for each cusp and one relation $\pm\mu_{j_1} \pm\mu_{j_2} \pm\mu_{j_3} = 0$ for each pair of pants. Cusps and pairs of pants correspond to faces and edges of P, and this gives $\mathrm{Ker}\,(p_*) = H^2(P; \mathbb{Z})$. □

Dehn filling Let us fix a triangulation T with n tetrahedra. By a Dehn filling of $D(T)$ we mean the result of some Dehn filling on some (possibly all) cusps of $D(T)$. We begin with the following observation, needed below:

Remark 3.6. If T is an ideal triangulation of a compact manifold $M(T)$ then the Dehn filling of $D(T)$ along all the longitudes is the orientable double $D(M(T))$ of $M(T)$. Moreover the cores of the filling solid tori are the doubles of the edges of T.

Proposition 3.5 now implies the following:

Corollary 3.7. *If in T there are at most n edges then every Dehn filling of $D(T)$ has positive first Betti number. In particular, it contains an incompressible surface.*

Proof. The assumption means that $D(T)$ has at most n cusps. A Dehn filling can decrease the first Betti number by at most one, whence the result. $\square$

Recall that the meridians of $D(T)$ were defined above. We now call *non-meridinal* a Dehn filling when it is non-meridinal on all cusps. The following is a consequence of the results of Agol and Lackenby on Dehn fillings.

Proposition 3.8. *Let V be a non-meridinal Dehn filling of $D(T)$.*

1. *If every edge of T has valence at least 6 then V is Haken and the core of each filling solid torus has infinite order in $\pi_1(V)$.*

2. *If every edge of T has valence at least 7 then V is hyperbolic.*

Proof. By Corollary 2.2, a non-meridinal slope on a cusp corresponding to an edge with valence q has length at least q in the boundary of the horospherical cusp neighbourhood $\mathcal{O}(T)$. Assume now all valences are at least 6. Then the filled loops have length at least 6 in $\mathcal{O}(T)$, so the proof of [15, Theorem 3.1] implies that V is irreducible and ∂-irreducible, and the cores of the filling tori have infinite order in $\pi_1(V)$. Since a tetrahedron has 6 edges, there can be at most n edges in T, so V is Haken by Corollary 3.7. This proves the first assertion. The second one now follows from the more general 6-theorem of Agol and Lackenby, together with Thurston's hyperbolization theorem for Haken manifolds. $\square$

Remark 3.9. The set $\mathcal{D}$ contains infinitely many examples, including Agol's one [1], which prove that the Agol-Lackenby theorem is sharp. To see this, take T with all edges having valence 6, and assume it is an ideal triangulation of an orientable manifold $M(T)$. As noted in Remark 3.6, the filling of $D(T)$ along the longitudes, which have length 6, is $D(M(T))$. But $D(M(T))$ is not hyperbolic, because $M(T)$ is hyperbolic with cusps (a structure is obtained by giving each tetrahedron of T the shape of a regular ideal one), so $D(M(T))$ contains essential tori.

Proposition 3.10. *Suppose that T is an ideal triangulation of a manifold $M(T)$ and that all the edges of T have valence at least 6. Then M is hyperbolic and no edge of T is homotopic with fixed ends to an arc contained in ∂M. In particular, the edges of T are properly homotopic to geodesics.*

Proof. We assign angles to the edges of the tetrahedra of T, choosing the angle $2\pi/v$ for an edge whose valence in T is v. The sum of angles around an edge of T is then of course always 2π, and the assumption implies that the sum of the three angles at a vertex of a tetrahedron is always at most π. This choice of angles gives T a structure weaker than the structure of *angled triangulation* defined by Lackenby [15], because he requires the sums of angles at the vertices to be precisely π. However one can check that the arguments in the proofs of [15, Propositions 4.4 and 4.5] extend *verbatim* to our triangulation with angles, and they imply that $M(T)$ contains no essential surface with non-negative Euler characteristic, so it is hyperbolic. By Remark 3.6 and Proposition 3.8, the doubles in $D(M(T))$ of the edges of T are homotopically non-trivial, so the edges of T cannot be homotopic to arcs in the boundary. $\square$

A closed incompressible surface An interesting property of $\mathcal{D}$ is described in the following:

Theorem 3.11. *Let T be a triangulation which consists of at least two tetrahedra. Then $D(T)$ contains a closed incompressible surface of genus 2 in its interior.*

Proof. Let F be a face adjacent to two distinct tetrahedra Δ_1 and Δ_2 of T. Let Σ_1 and Σ_2 be two (non-geodesic) parallel copies of the thrice-punctured sphere in $D(T)$ corresponding to F, each Σ_i contained inside the block K corresponding to Δ_i. The surfaces Σ_1 and Σ_2 cut the cusps they are incident to into 6 annuli, three small ones between them and three other ones A_1, A_2, and A_3. We then define Σ as $\Sigma_1 \cup \Sigma_2 \cup A_1 \cup A_2 \cup A_3$. Since Σ can be pushed in the interior of $D(T)$, it is sufficient to show that it is incompressible.

Suppose by contradiction that Σ is compressible in $D(T)$. Let us choose among all the compressing discs for Σ which are transverse to the thrice-punctured spheres S_i corresponding to the faces of T the one, say D, whose intersection with $\bigcup S_i$ consists of the minimal number k of components. We claim that $k = 0$. If $D \cap (\bigcup S_i)$ has a component which is a loop then the loop must be trivial in one of the S_i's (because they are geodesic surfaces). An easy innermost argument then contradicts the minimality of k.

Suppose then that $D \cap (\bigcup S_i)$ has arc components, and choose an outermost one, so there is an arc $\beta \subset \partial D$ such that $\alpha \cup \beta = \partial D'$, with D' a disc contained in D. Note that D' is embedded in a block K as shown in Fig. 20, with α lying on one S_i and β on Σ. By construction, the surface Σ is cut by the S_i's into some tubes and two pair of pants P_1 and P_2, contained in distinct K's. Each such tube or pair of pants is isotopic to a tube in a cusp or a pair of pants inside a S_i, so it is incompressible.

The arc β lies in a tube or in a pair of pants P_i. If β is not essential in it, then D' can be isotoped away, against the minimality of k. This is always the case when β lies on a tube, because $\alpha \subset S_i$ and $\alpha \cup \beta$ is a loop, so the ends of β must be on the same boundary component. Therefore β lies on a P_i, and again the ends of β must be on the same boundary component of the tube. It easily follows that $\alpha \cup \beta$ is not null-homologous in K, but $\alpha \cup \beta = \partial D'$, whence a contradiction. This implies that indeed $k = 0$.

We have shown that if Σ is compressible then the intersection of Σ with some block K is compressible within the block, but in the course of our argument we have also noticed that this cannot be the case. We then get a contradiction, and the proof is complete. $\square$

4 Isometry groups

Kojima proved in [13] that any finite group is the full isometry group of a closed hyperbolic 3-manifold. We give here a new proof of this result, also providing an upper bound on the volume of the manifold in terms of the order of the group:

Theorem 4.1. *There exists $k > 0$ such that the following holds: for any finite group G there exists a closed orientable hyperbolic 3-manifold V_G with* $\mathrm{Isom}(V_G) = \mathrm{Isom}^+(V_G) \cong G$ *and* $\mathrm{Vol}(V_G) \leqslant k \cdot |G|^9$.

Before plunging into the details of the proof we briefly describe the scheme of our argument. In the whole section we fix a finite group G and we denote by n its order. We first construct a special polyhedron Q_G with $\pi_1(Q_G) \cong G$. We carefully choose a Q_G having no symmetries, in such a way that the group of combinatorial automorphisms of its universal cover P_G does not exceed the group of deck transformations, which is isomorphic to G. Now let T_G be the triangulation dual to P_G. Corollary 2.5 ensures that $\mathrm{Isom}^+(D(T_G)) \cong G$. In order to kill the orientation-reversing isometries of $D(T_G)$ we then perform on $D(T_G)$ a suitable Dehn filling, which finally gives the desired V_G. Since volume decreases under Dehn filling and the volume of $D(T_G)$ is proportional to the number $c(P_G)$ of vertices of P_G, the bound on the volume follows from a bound on $c(P_G)$.

Construction of the polyhedron We begin with the following:

Lemma 4.2. *There exists $k > 0$ independent of G and a special polyhedron Q_G with $\pi_1(Q_G) \cong G$ and $c(Q_G) \leqslant k \cdot n^4$.*

Proof. For the trivial group we can take a special spine of S^3, so we assume $n \geqslant 2$. We consider the trivial presentation for G having as generators all the elements of $G \setminus \{1\}$ and as relations all the expressions abc^{-1} if $a, b \neq 1$ and $a \cdot b = c \neq 1$ in G, and all the expressions ab if $a \neq 1$ and $a \cdot b = 1$ in G. Let U_n be the polyhedron obtained by adding $n-1$ handles to S^2, thus getting a surface Σ_n, and then attaching one disc to the core of each handle. We now fix a bijective correspondence between $G \setminus \{1\}$ and the handles, getting an isomorphism between the free group on $G \setminus \{1\}$ and $\pi_1(U_n)$.

The relations in the presentation of G translate into simple loops $\gamma_1, \ldots, \gamma_{(n-1)^2}$ on Σ_n. We choose the γ_i's in generic position with respect to each other and to the cores of the handles, and intersecting each other and the cores in a minimal number ν of points. The first condition and the fact that $n \geqslant 2$ easily imply that by attaching discs to U_n along the γ_i's we get a special polyhedron Q_G with $\pi_1(Q_G) \cong G$ and ν vertices. The γ_i's run $3n-4$ times (with multiplicity) along each handle of Σ_n, and they give rise to $(n-1)(3n-4)$ properly embedded arcs in the central punctured sphere of U_n. So $(n-1)(3n-4)$ vertices of Q_G arise along the handles of Σ_n. The other vertices lie in the central sphere of U_n, and there can be at most $(n-1)(3n-4)(3n^2-7n+3)/2$ of them. Therefore $\nu \leqslant k \cdot n^4$ for some $k > 0$ independent of G. $\square$

Proposition 4.3. *There exists $k > 0$ independent of G and a special polyhedron P_G with group of combinatorial automorphisms isomorphic to G and $c(P_G) \leqslant k \cdot n^9$.*

Proof. We suitably modify the polyhedron Q_G given by the previous lemma to ensure that its universal cover has the desired properties.

We begin with some definitions. If P is a special polyhedron, we say that a vertex v of P is *good* if there is no edge of P with both ends at v. If v is a good vertex of P and e is an edge incident to v, we define a non-negative integer $\lambda(v, e)$ as the maximal length of a simple simplicial path in the singular set of P starting at v with direction e and touching, apart from v, bad vertices only.

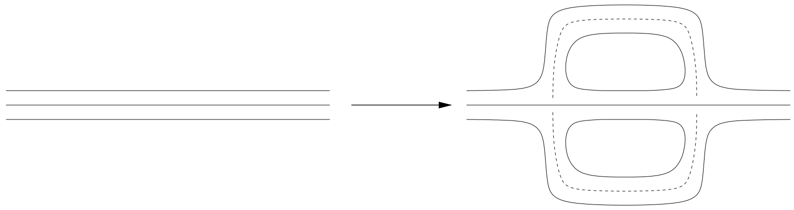

Figure 21: Performing this move along each edge we eliminate bad vertices. The picture shows a neighbourhood of the singular set, which determines a special polyhedron uniquely.

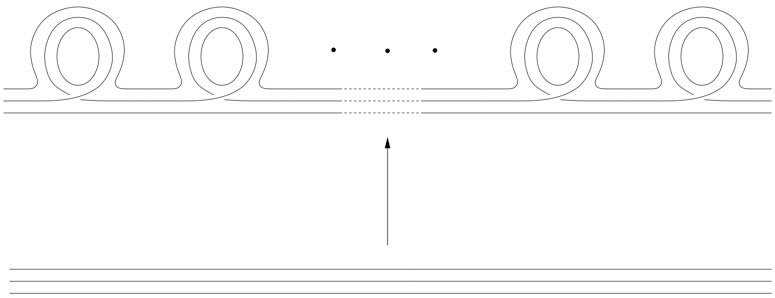

Figure 22: Adding some curls along an edge of a special polyhedron.

We first modify Q_G by performing the local move described in Fig. 21 along each edge of Q_G. In this way we obtain a polyhedron Q'_G without bad vertices and such that $\pi_1(Q'_G) \cong \pi_1(Q_G)$ and $c(Q'_G) = 5 \cdot c(Q_G)$. We now call *curl* in a special polyhedron P a face of P whose boundary consists of a single (closed) edge of P. Let $(e_1, \ldots, e_\nu)$ be any ordering of the edges of Q'_G, so $\nu = 2 \cdot c(Q'_G)$, and modify Q'_G by adding i curls along e_i as described in Fig. 22, thus obtaining a polyhedron Q''_G. Observe that Q''_G is homotopically equivalent to Q'_G and that $c(Q''_G) = 2c(Q'_G)(c(Q'_G)+1)$. We now define P_G to be the universal cover of Q''_G. The upper bound on $c(P_G)$ being obvious, we are left to prove the first required property.

Let $\pi : P_G \to Q''_G$ be a fixed covering projection. We first observe that the bad vertices of Q''_G are precisely the $c(Q'_G)(2c(Q'_G)+1)$ vertices added to Q'_G. Moreover, any edge f starting from a good vertex v of Q''_G comes from a unique edge $e_{i(f)}$ of Q'_G. By construction we have $i(f) = \#\{\text{curls along } e_{i(f)}\} = \lambda(v, f)$. The curls added to Q'_G to get Q''_G are contractible, so they isomorphically lift to P_G. This easily implies that a vertex of P_G is good if and only if its projection in Q''_G is. Moreover, if $\widetilde{v} \in \pi^{-1}(v)$ is a good vertex of P_G and $\widetilde{f}$ is an edge starting from $\widetilde{v}$ with $\pi(\widetilde{f}) = f$, then $\lambda(\widetilde{v}, \widetilde{f}) = \lambda(v, f) = i(f)$. Now let φ be a combinatorial automorphism of P_G, let $\widetilde{v}$ be a good vertex of P_G and let $\widetilde{f}_1, \widetilde{f}_2, \widetilde{f}_3, \widetilde{f}_4$ be the edges emanating from $\widetilde{v}$. Of course we have $\lambda(\varphi(\widetilde{v}), \varphi(\widetilde{f}_i)) = \lambda(\widetilde{v}, \widetilde{f}_i)$ for $i = 1, 2, 3, 4$. Since Q'_G has no bad vertices and at least three good vertices, this forces $\pi(\varphi(\widetilde{v})) = \pi(\widetilde{v})$ and $\pi(\varphi(\widetilde{f}_i)) = \pi(\widetilde{f}_i)$ for $i = 1, 2, 3, 4$. It easily follows that $\pi \circ \varphi = \pi$, *i.e.* that φ is a deck transformation. $\square$

Dehn filling Let T_G be the triangulation dual to the special polyhedron P_G constructed during the proofs of Lemma 4.2 and Proposition 4.3. The next result concludes the proof of Theorem 4.1.

Proposition 4.4. *There exists a hyperbolic Dehn filling V_G of $D(T_G)$ with* $\mathrm{Isom}(V_G) = \mathrm{Isom}^+(V_G) \cong G$.

Proof. We denote $N(T_G)$ by N and $D(T_G)$ by D. Since T_G has at least three good vertices, it is not one of the exceptional triangulations defined in Fig. 8. Then Corollary 2.5 implies that $\mathrm{Isom}^+(D) \cong G$ and that there is a canonical orientation-reversing involution τ of D. Let $C_1, \ldots, C_m$ be the cusps of D, and let E_i be the boundary torus corresponding to C_i of the compactification of D. Note that τ extends to the E_i's. Since τ leaves invariant each C_i and reverses the orientation on it, there exists exactly one slope s_i on E_i with $\tau(s_i) = s_i$.

Since $\mathrm{Aut}(P_G)$ is a group of covering transformations and the 2-dimensional regions of P_G are discs, Brouwer's fixed point theorem implies that $\mathrm{Aut}(P_G)$ acts freely on the regions, so $\mathrm{Isom}^+(D) \cong \mathrm{Isom}(N) \cong \mathrm{Aut}(P_G)$ acts freely on the cusps. This implies that there exist systems $(l_1, \ldots, l_m)$ of one slope per cusp which are invariant under the action of $\mathrm{Isom}^+(D)$. More precisely, Thurston's hyperbolic Dehn filling theorem implies that there exists $(l_1, \ldots, l_m)$ with the following properties:

(A) $(l_1, \ldots, l_m)$ is $\mathrm{Isom}^+(D)$-invariant;

(B) the filling $D(l_1, \ldots, l_m)$ of D along the l_i's is hyperbolic;

(C) the cores of the filling tori are the m shortest geodesics of $D(l_1, \ldots, l_m)$;

(D) $l_i \neq s_i$ for $i = 1, \ldots, m$.

If h is an element of $\mathrm{Isom}^+(D)$, from (A) we see that h extends to an automorphism of $D(l_1, \ldots, l_m)$, which in turn is homotopic by Mostow's rigidity to an isometry $e(h)$. It is easily seen that the map $e : \mathrm{Isom}^+(D) \to \mathrm{Isom}(D(l_1, \ldots, l_m))$ thus defined is an injective group homomorphism. To show it is surjective, pick $g \in \mathrm{Isom}(D(l_1, \ldots, l_m))$. By (C), the cores of the filling tori are g-invariant, so g restricts to an automorphism of D, which by rigidity is homotopic to an isometry $r(g)$. Moreover $r(g)$ leaves $(l_1, \ldots, l_m)$ invariant. Now $r(g) = \tau^\varepsilon \circ h$ for some $h \in \mathrm{Isom}^+(D)$ and $\varepsilon \in \{0, 1\}$. If $r(g)$ maps C_1 to C_i then $r(g)(l_1) = l_i$, but we know $h(l_1) = l_i$ and $\tau(l_i) \neq l_i$ by (D), so $\varepsilon = 0$. It follows that $r(g) = h \in \mathrm{Isom}^+(D)$, so $g = e(h)$. □

References

[1] I. Agol, *Bounds on exceptional Dehn filling*, Geom. Topol. **4** (2000), 431-449.

[2] ———, *Small* 3-*manifolds of large genus*, Geom. Dedicata **102** (2003), 53-64.

[3] G. Amendola, *A calculus for ideal triangulations of three-manifolds with embedded arcs*, `math.GT/0301219`, to appear in Math. Nachr.

[4] F. Costantino, D. Thurston, *Shadow complexity of* 3-*manifolds*, in preparation.

[5] C. J. Earle, J. Eells, *A fibre bundle description of Teichmüller theory*, J. Differential Geometry **3** (1969), 19-43.

[6] D. B. A. Epstein, *Curves on* 2-*manifolds and isotopies*, Acta Math. **115** (1966), 83-107.

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

[8] R. Frigerio – B. Martelli – C. Petronio, *Complexity and Heegaard genus of an infinite class of hyperbolic* 3-*manifolds*, Pacific J. Math. **210** (2003), 283-297.

[9] ———, *Small hyperbolic 3-manifolds with geodesic boundary*, `math.GT/0211425`.

[10] R. Frigerio – C. Petronio, *Construction and recognition of hyperbolic* 3-*manifolds with geodesic boundary*, `math.GT/0109012`, to appear in Trans. Amer. Math. Soc.

[11] C. McA. Gordon, *Small surfaces and Dehn filling*, in "Proceedings of the Kirbyfest" (Berkeley, CA, 1998), pp. 177-199 (electronic), Geometry and Topology Monographs, Vol. 2, Coventry, 1999.

[12] K. Johannson, "Homotopy equivalences of 3-manifolds with boundaries," Lecture Notes in Mathematics, Vol. 761, Springer-Verlag, Berlin, 1979.

[13] S. Kojima, *Isometry tranformations of hyperbolic 3-manifolds*, Topology Appl. **29** (1988), 297-307.

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

[15] M. Lackenby, *Word hyperbolic Dehn surgery*, Invent. Math. **140** (2000), 243-282.

[16] Y. Miyamoto, *Volumes of hyperbolic manifolds with geodesic boundary*, Topology **33** (1994), 613-629.

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

[18] V. G. Turaev, "Quantum invariants of knots and 3-manifolds," de Gruyter Studies in Mathematics, Vol. 18, Walter de Gruyter & Co., Berlin, 1994.

[19] F. Waldhausen, *On irreducible 3-manifolds which are sufficiently large*, Ann. of Math. (2) **87** (1968), 56-88.

[20] J. R. Weeks, *SnapPea*, The hyperbolic structures computer program, available from `www.geometrygames.org`.

Scuola Normale Superiore
Piazza dei Cavalieri 7, 56127 Pisa, Italy
f.costantino@sns.it

Dipartimento di Matematica, Università di Pisa
Via F. Buonarroti 2, 56127 Pisa, Italy
frigerio@mail.dm.unipi.it, martelli@mail.dm.unipi.it

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