arXiv:math.GT/0409606 v1 30 Sep 2004

Spherical splitting of 3-orbifolds

Carlo Petronio

October 1, 2004

Abstract

The famous Haken-Kneser-Milnor theorem states that every 3-manifold can be expressed in a unique way as a connected sum of prime 3-manifolds. The analogous statement for 3-orbifolds has been part of the folklore for several years, and it was commonly believed that slight variations on the argument used for manifolds would be sufficient to establish it. We demonstrate in this paper that this is not the case, proving that the apparently natural notion of "essential" system of spherical 2-orbifolds is not adequate in this context. We also show that the statement itself of the theorem must be given in a substantially different way. We then prove the theorem in full detail, using a certain notion of "efficient splitting system."

MSC (2000): 57M99 (primary), 57M50, 57M60 (secondary).

Introduction

Since the seminal work of Thurston [10], the concept of orbifold has become a central one in 3-dimensional geometric topology, because it embodies the notion of a 3-manifold and that of a finite group action on that 3-manifold. As a matter of fact, Thurston himself indicated how to prove that, under some natural topological constraints, a 3-orbifold with non-empty singular locus carries one of the 8 geometries, so it is a quotient of a geometric 3-manifold under a finite isometric action (see [1] and [4] for modern detailed accounts of the argument).

In the case of manifolds (closed orientable ones, say) the two basic obstructions to the existence of a geometric structure are related to the existence of essential spheres and of essential tori. So one restricts first to irreducible manifolds and next to Seifert and atoroidal ones. For this restricted class of manifolds then one has Thurston's *geometrization conjecture* (or *theorem*, according to the recent work of Perelman [8]). This result can be regarded as a *uniformization theorem* for 3-manifolds in view of the Haken-Kneser-Milnor theorem (canonical splitting of a 3-manifold along spheres into irreducible 3-manifolds) and the Jaco-Shalen-Johansson theorem (canonical splitting of an irreducible 3-manifold along tori into Seifert and atoroidal 3-manifolds).

The case of 3-orbifolds is similar, where the first basic obstruction to the existence of a geometric structure is removed by restricting to (suitably defined) *irreducible* 3-orbifolds. The natural expectation would then be to have an analogue of the Haken-Kneser-Milnor theorem, asserting that each 3-orbifold canonically splits along spherical 2-orbifolds into irreducible 3-orbifolds. This result has actually been part of the folklore for several years, but it turns out that to state and prove it in detail one has to somewhat refine the ideas underlying the proof in the manifold case. This paper is devoted to these refinements.

We will state in the present introduction two versions of the splitting theorem that we will establish below, addressing the reader to Section 1 for the appropriate definitions.

Theorem 0.1. *Let X be a closed locally orientable 3-orbifold. Suppose that X does not contain any bad 2-suborbifold, and that every spherical 2-suborbifold of X is separating. Then X contains a finite system $\mathcal{S}$ of spherical 2-suborbifolds such that:*

- *No component of $X \setminus \mathcal{S}$ is punctured discal.*
- *If Y is the orbifold obtained from $X \setminus \mathcal{S}$ by capping the boundary components with discal 3-orbifolds, then the components of Y are irreducible.*

Any such system $\mathcal{S}$ and the corresponding Y, if viewed as abstract collections of 2- and 3-orbifolds respectively, depend on X only.

We inform the reader that this result has a dual interpretation in terms of connected sum, discussed in detail in Section 1. We also note that the existence part of Theorem 0.1 may seem to have a purely theoretical nature, which is not pleasant in the context of 3-dimensional topology, where a great emphasis is typically given to algorithmic methods [7]. The next longer statement emphasizes the constructive aspects of the splitting.

Theorem 0.2. *Let X be a closed locally orientable 3-orbifold. Suppose that X does not contain any bad 2-suborbifold, and that every spherical 2-suborbifold of X is separating. Starting with $\mathcal{S} = \emptyset$ and $Y = X$, consider the (non-deterministic) process described by the following steps:*

1. *If all ordinary spherical* 2*-suborbifolds of* Y *are inessential, turn to Step 2. Otherwise choose* Σ *as one such* 2*-suborbifold, redefine* $\mathcal{S}$ *as the given* $\mathcal{S}$ *union* Σ*, redefine* Y *as the orbifold obtained from the given* Y *by cutting along* Σ *and capping the result by discal* 3*-orbifolds, and repeat Step 1.*

2. *If all cyclic spherical* 2*-suborbifolds of* Y *are inessential, turn to Step 3. Otherwise choose* Σ *as one such* 2*-suborbifold, redefine* $\mathcal{S}$ *and* Y *as in Step 1, and repeat Step 2.*

3. *If all vertex spherical* 2*-suborbifolds of* Y *are inessential, output* $\mathcal{S}$ *and* Y*. Otherwise choose* Σ *as one such* 2*-suborbifold, redefine* $\mathcal{S}$ *and* Y *as in Step 1, and repeat Step 3.*

Then the process is finite and the final Y *consists of irreducible* 3*-orbifolds. Moreover, the final* $\mathcal{S}$ *and* Y*, viewed as abstract collections of* 2*- and* 3*-orbifolds respectively, depend on* X *only, not on the specific process leading to them. In addition, at all steps of the process the essential spherical* 2*-orbifold* Σ *can be chosen to be normal with respect to a triangulation of the given orbifold* Y*.*

For the reader familiar with the proof of the splitting theorem for manifolds (see for instance [3]) we will now point out what goes wrong in the orbifold case, and sketch how we have modified the argument. The main steps to show the existence of a splitting of a manifold M into prime ones are as follows:

(a) Show that if M contains a non-separating sphere then it splits as the connected sum of $S^2 \times S^1$ or $S^2 \mathbin{\tilde{\times}} S^1$ and some M', and the first Betti number of M' is smaller than that of M. Deduce that one can restrict to M's in which all spheres are separating.

(b) Define a family of spheres $\mathcal{S} \subset M$ to be *essential* if no component of $M \setminus \mathcal{S}$ is a punctured disc. Show that every essential $\mathcal{S}$ can be replaced by one having the same number of components and being normal with respect to any given triangulation of M. Deduce that all maximal essential systems of spheres are finite.

(c) Prove that if $\mathcal{S}$ is a maximal essential system then the manifolds obtained by capping off the components of $M \setminus \mathcal{S}$ are irreducible.

Uniqueness is more elaborate to prove, but its validity only depends on the following fact, that we express using the dual viewpoint of connected sum:

(d) The sphere S^3 is the only identity element of the operation of connected sum. No connected sum with S^3 should be employed to realize a manifold.

Turning to orbifolds, we can show that none of the points (a)-(d) has a straight-forward extension:

($\neg$a) The existence of a non-separating spherical 2-suborbifold does not imply the existence of an essential separating one (see Fig.3 below). In particular, there are infinitely many 3-orbifolds containing some non-separating but no essential separating spherical 2-orbifold (see Section 5).

The difficulty given by non-separating spherical 2-orbifolds is mentioned in [2], where the choice is made of cutting along these 2-orbifolds too. The reasons of our choice of excluding non-separating spherical 2-orbifolds *tout-court* will be discussed in Section 3. We only want to mention here that the techniques of the present paper do not seem to adapt easily to include the case of splitting along non-separating spherical 2-orbifolds.

Turning to point (b), we follow [2] and define a family $\mathcal{S}$ of spherical 2-suborbifolds of a 3-orbifold X to be *essential* if no component of $X \setminus \mathcal{S}$ is a punctured discal 3-orbifold. We now have:

($\neg$b) If $\mathcal{S}$ is essential and $\mathcal{S}', \mathcal{S}''$ are obtained from $\mathcal{S}$ by one of the compression moves necessary to normalize $\mathcal{S}$, then neither of $\mathcal{S}'$ and $\mathcal{S}''$ need be essential (see Fig. 4 below).

This fact could be viewed as a merely technical difficulty, showing that the theory of normal surfaces is inadequate to prove finiteness of maximal essential systems. The following fact is however more striking:

($\neg$c) If $\mathcal{S}$ is a maximal essential system of spherical 2-suborbifolds of X, and Y is obtained by cutting X along $\mathcal{S}$ and capping the result, the components of Y need not be irreducible (see Fig. 5 below).

Given this drawback of the notion of essential system, we have refined it to a certain notion of *efficient splitting system*, where we impose artificially that the capped components of the complement should be irreducible. The proof of the existence of efficient splitting systems then requires some more efforts than in the manifold case. In addition, finiteness of the splitting process in Theorem 0.2 does not follow directly from the existence of finite efficient splitting systems.

Concerning uniqueness, we have the following:

(¬d) There are three different types of connected sums of orbifolds, and each has its own identity, different from the other ones. In a sequence of connected sums of irreducible 3-orbifolds, the property that a certain sum is "trivial" depends on the order in which the sums are performed.

The same notion of efficient system used to prove the existence of a splitting actually allows to deal also with this problem, and hence to prove uniqueness in Theorem 0.1. The same argument employed for uniqueness also proves finiteness of the splitting process in Theorem 0.2.

Acknowledgements The author warmly thanks Gwénaël Massuyeau, Sergei Matveev, Luisa Paoluzzi, and Joan Porti for helpful discussions.

1 Efficient connected sums and splitting systems

In this section we introduce the necessary terminology and we provide alternative statements of Theorems 0.1 and 0.2.

Local structure of orbifolds We will not cover here the general theory of orbifolds, referring the reader to the milestone [10], to the excellent and very recent [2], and to the comprehensive bibliography of the latter. We just recall that an orbifold of dimension n is a topological space with a singular smooth structure, locally modelled on a quotient of $\mathbb{R}^n$ under the action of a finite group of diffeomorphisms. We will only need to refer to the cases $n = 2$ and $n = 3$, and we will confine ourselves to orientation-preserving diffeomorphisms. In addition, all our orbifolds will be compact.

Under these assumptions one can see that a 2-orbifold Σ consists of a compact *support* surface $|\Sigma|$ together with a finite collection $S(\Sigma)$ of points in the interior of $|\Sigma|$, the *cone* points, each carrying a certain *order* in $\{p \in \mathbb{N} : \ p \geqslant 2\}$.

Analogously, a 3-orbifold X is given by a compact *support* 3-manifold $|X|$ together with a *singular set* $S(X)$. Here $S(X)$ is a finite collection of circles and unitrivalent graphs tamely embedded in $|X|$, with the univalent vertices given by the intersection with $\partial|X|$. Moreover each component of $S(X)$ minus the vertices carries an *order* in $\{p \in \mathbb{N} : \ p \geqslant 2\}$, with the restriction that the three germs of edges incident to each vertex should have orders $(2, 2, p)$, for arbitrary p, or $(2, 3, p)$, for $p \in \{3, 4, 5\}$.

Bad, spherical, and discal orbifolds An *orbifold-covering* is a map between orbifolds locally modelled on a map of the form $\mathbb{R}^n/_\Delta \to \mathbb{R}^n/_\Gamma$, naturally defined whenever $\Delta < \Gamma < \mathrm{Diff}_+(\mathbb{R}^n)$. An orbifold is called *good* when it is orbifold-covered by a manifold, and *bad* when it is not good. And we call it *spherical* (respectively, discal) when it is the quotient of the sphere $S^n = \mathbb{R}^n \cup \{\infty\}$ (respectively, the disc D^n) under the action of a finite group of $\mathrm{SO}(n)$. In the sequel we will need the following easy results:

Lemma 1.1. *The only bad closed 2-orbifolds are $(S^2; p)$, the 2-sphere with one cone point of order p, and $(S^2; p, q)$, the 2-sphere with cone points of orders $p \neq q$.*

To state the other results, we introduce some notation repeatedly used below. We define $\mathbb{D}^3_{\mathrm{o}}$ to be D^3, the *ordinary* 3-disc, $\mathbb{D}^3_{\mathrm{c}}(p)$ to be D^3 with singular set a trivially embedded arc with arbitrary order p, and $\mathbb{D}^3_{\mathrm{v}}(p,q,r)$ to be D^3 with singular set a trivially embedded "Y-graph" with edges of orders p, q, r. We will call $\mathbb{D}^3_{\mathrm{c}}(p)$ and $\mathbb{D}^3_{\mathrm{v}}(p,q,r)$ respectively *cyclic* and *vertex 3-discs*, and we will employ the shortened notation $\mathbb{D}^3_{\mathrm{c}}$ and $\mathbb{D}^3_{\mathrm{v}}$ to denote cyclic and vertex 3-discs with generic orders. We now define the *ordinary, cyclic*, and *vertex 2-spheres*, denoted respectively by $\mathbb{S}^2_{\mathrm{o}}$, $\mathbb{S}^2_{\mathrm{c}}(p)$, and $\mathbb{S}^2_{\mathrm{v}}(p,q,r)$, as the 2-orbifolds bounding the corresponding 3-discs $\mathbb{D}^3_{\mathrm{o}}$, $\mathbb{D}^3_{\mathrm{c}}(p)$ and $\mathbb{D}^3_{\mathrm{v}}(p,q,r)$. We also define the *ordinary, cyclic*, and *vertex 3-spheres*, denoted respectively by $\mathbb{S}^3_{\mathrm{o}}$, $\mathbb{S}^3_{\mathrm{c}}(p)$, and $\mathbb{S}^3_{\mathrm{v}}(p,q,r)$, as the 3-orbifolds obtained by mirroring the corresponding 3-discs $\mathbb{D}^3_{\mathrm{o}}$, $\mathbb{D}^3_{\mathrm{c}}(p)$, and $\mathbb{D}^3_{\mathrm{v}}(p,q,r)$ in their boundary. The 2- and 3-spheres with generic orders will be denoted by $\mathbb{S}^2_*$ and $\mathbb{S}^3_*$.

Lemma 1.2. *The spherical 2-orbifolds are precisely the $\mathbb{S}^2_*$'s.*

Lemma 1.3. *The discal 3-orbifolds are precisely the $\mathbb{D}^3_*$'s.*

Lemma 1.4. *The spherical 3-orbifolds are precisely the $\mathbb{S}^3_*$'s.*

2-suborbifolds and irreducible 3-orbifolds We say that a 2-orbifold Σ is a *suborbifold* of a 3-orbifold X if $|\Sigma|$ is embedded in $|X|$ so that $|\Sigma|$ meets $S(X)$ transversely (in particular, it does not meet the vertices), and $S(\Sigma)$ is given precisely by $|\Sigma| \cap S(X)$, with matching orders.

A spherical 2-suborbifold Σ of a 3-orbifold X is called *essential* if it does not bound in X a discal 3-orbifold. A 3-orbifold is called *irreducible* if it does not contain any bad 2-suborbifold and every spherical 2-suborbifold is inessential (in particular, it is separating).

If a 3-orbifold X is bounded by spherical 2-orbifolds, we can canonically associate to X a closed orbifold $\widehat{X}$ by attaching the appropriate discal 3-orbifold to each component of ∂X. We say that $\widehat{X}$ is obtained by *capping* X.

Whenever a 3-orbifold is not irreducible, we can select an essential spherical 2-suborbifold Σ, split X along Σ, and cap the result. A naïf statement of the theorem we want to prove would be that *the successive process of splitting and capping comes to an end in finite time, and the result is a collection of irreducible* 3*-orbifolds, which depends on* X *only.* With the appropriate assumptions and choices of the essential spherical 2-orbifolds this is eventually true (Theorem 0.2) but, just as in the case of manifolds, the statement and proof are best understood by performing a simultaneous splitting along a system of spherical 2-orbifolds, rather than a successive splitting along single ones.

Connected sum To motivate the statement of the main result given below, we reverse our point of view. The remark here is that, if X can be turned into a collection of irreducible 3-orbifolds by the "split and cap" strategy mentioned above, then X can be reconstructed from the same collection by "puncture and glue" operations. As in the case of manifolds, each such operation will be called a *connected sum*. But the world of orbifolds is somewhat more complicated than that of manifolds, so we need to give the definition with a little care.

Let X_1 and X_2 be 3-orbifolds, and pick points $x_i \in |X_i|$, such that one of the following holds:

(i) Both x_1 and x_2 are non-singular points.

(ii) Both x_1 and x_2 are singular but not vertex points, and they belong to singular edges of the same order p.

(iii) Both x_1 and x_2 are vertex singular points, and the triple p, q, r of orders of the incident edges is the same for x_1 and x_2.

We can then remove from $|X_i|$ a regular neighbourhood of x_i, and glue together the resulting orbifolds by a homeomorphism which matches the singular loci. The result is a 3-orbifold X that we call a *connected sum* of X_1 and X_2. Depending on which of the conditions (i), (ii), or (iii) is met, we call the operation (and its result), a connected sum of *ordinary* type, of *cyclic* type *(of order* p*)*, or of *vertex* type *(of order* p, q, r*)*.

Trivial connected sums In the case of manifolds, there is only one type of operation of connected sum, and its only identity element is the ordinary 3-sphere. For 3-orbifolds we have the easy:

Remark 1.5. The unique identity element of the ordinary connected sum (respectively, of the order-p cyclic connected sum, and of the the order-(p,q,r) vertex connected sum) is $\mathbb{S}^3_{\mathrm{o}}$ (respectively, $\mathbb{S}^3_{\mathrm{c}}(p)$, and $\mathbb{S}^3_{\mathrm{v}}(p,q,r)$).

We can then define to be *trivial* an operation of connected sum with the corresponding identity element. However, when one considers a sequence of operations of connected sums, one must be careful, because a sum which appears to be non-trivial in the first place may actually turn out to be trivial later on, as the following example, suggested by Joan Porti, shows.

Example 1.6. Let K be any non-trivial knot in S^3, and let $p,q \geqslant 2$ be distinct integers. Construct a 3-orbifold X by first taking the ordinary connected sum of $\mathbb{S}^3_{\mathrm{c}}(p)$ and $\mathbb{S}^3_{\mathrm{c}}(q)$, and then taking the order-p cyclic connected sum of the result with (S^3, K_p), the orbifold with support S^3 and singular set K of order p. Then both sums are non-trivial. However, if we reverse the order, we can realize the same X as the order-p cyclic sum of $\mathbb{S}^3_{\mathrm{c}}(p)$ and (S^3, K_p), followed by the ordinary sum with $\mathbb{S}^3_{\mathrm{c}}(q)$, and the first sum is now trivial.

This phenomenon is of course a serious obstacle to proving that each 3-orbifold has a unique expression as a connected sum of irreducible ones. To overcome this obstacle, we need to redefine the notion of trivial connected sum.

Graphs and efficient connected sums We consider in this paragraph successive connected sums $X_0\#\ldots\#X_n$, stipulating with this notation that the first sum is between X_1 and X_0, the second sum is between X_2 and $X_0\#X_1$, and so on. Note first that to perform the j-th sum we can always arrange the puncture made to $X_0\#\ldots\#X_{j-1}$, to be disjoint from the spherical 2-orbifolds along which the punctured $X_0,\ldots,X_{j-1}$ have been glued. This implies that X_j is glued to precisely one of $X_0,\ldots,X_{j-1}$. We can then construct a graph associated to the connected sum, with nodes the X_j's, and edges labelled by the types of sums performed. Note that this graph is actually a tree.

We will say that a connected sum $X_0\#\ldots\#X_n$ is *efficient* if the X_j's are irreducible and in its graph there is no node $\mathbb{S}^3_{\mathrm{o}}$ with an incident edge of ordinary type, there is no node $\mathbb{S}^3_{\mathrm{c}}$ with an incident edge of cyclic type, and there is no node $\mathbb{S}^3_{\mathrm{v}}$ with an incident edge of vertex type, see Fig. 1.

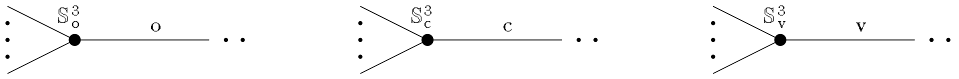

Figure 1: Inefficient connected sums.

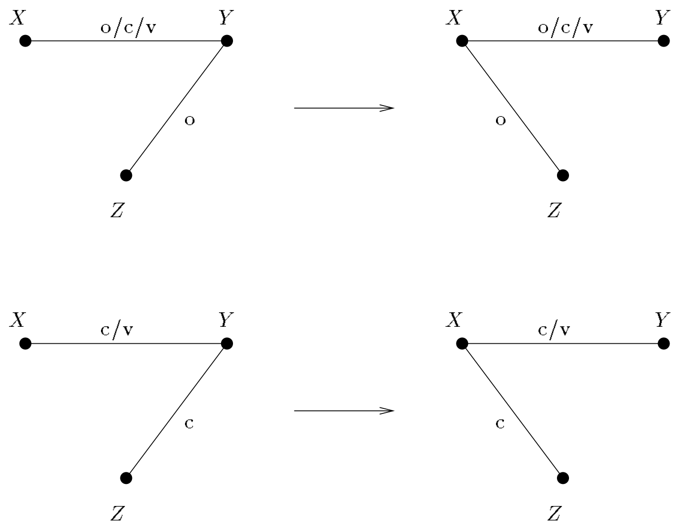

Figure 2: Moves relating two graphs of the same connected sum.

One should notice that the graph associated to $X_0\#\ldots\#X_n$ is not quite unique, since to construct it one needs to isotope each puncture away from the previous ones, as described above. However, one easily sees that two graphs of $X_0\#\ldots\#X_n$ differ by moves as in Fig. 2, and one deduces that the notion of efficient connected sum is indeed well-defined. The following is an alternative statement of our main result:

Theorem 1.7. *Let X be a closed locally orientable 3-orbifold. Suppose that X does not contain any bad 2-suborbifold, and that every spherical 2-suborbifold of X is separating. Then X can be realized as an efficient connected sum of irreducible 3-orbifolds. Any two realizations involve the same irreducible summands and the same types of sums (including orders).*

Efficient splitting systems As in the case of manifolds, to establish the existence part of Theorem 1.7 one proceeds with the original strategy of "split and cap" first mentioned above, but one has to split symultaneously along a (finite) system of disjoint spherical 2-orbifolds.

To this end we give the following definitions. We say that a 3-orbifold is *punctured discal* if it is obtained from some $\mathbb{D}^3_*$ by removing a regular neighbourhood of a finite set. A system $\mathcal{S}$ of spherical 2-suborbifolds of X is called *essential* if no component of $X \setminus \mathcal{S}$ is punctured discal, and it is called *coirreducible* if all the components of $(X \setminus \mathcal{S})^\frown$ are irreducible. We call *efficient splitting system* a finite system of spherical 2-suborbifolds of X which is both essential and coirreducible.

Remark 1.8. To each system $\mathcal{S}$ of separating spherical 2-suborbifolds there corresponds a realization of X as a connected sum of the components of $(X \setminus \mathcal{S})^\frown$. Moreover $\mathcal{S}$ is efficient if and only if the connected sum is.

The result we will establish directly below, which is obviously equivalent to Theorems 0.1 and 1.7, and will be used to deduce Theorem 0.2, is the following:

Theorem 1.9. *Let X be a closed locally orientable 3-orbifold. Suppose that X does not contain any bad 2-suborbifold, and that every spherical 2-suborbifold of X is separating. Then X admits efficient splitting systems. Any two splitting systems coincide as abstract collections of spherical 2-orbifolds, and the capped components of their complements also coincide.*

2 Comparison with the manifold case

We illustrate in greater detail in this section why it has been necessary to substantially modify the proof of the splitting theorem in passing from the manifold to the orbifold case.

Concerning the issue of non-separating spherical 2-orbifolds (point (a) of the Introduction) we begin by recalling what is the point for an orientable manifold M. If $\Sigma \subset M$ is a non-separating sphere, one chooses a simple closed curve α meeting Σ once and transversely and one takes the boundary Σ' of a regular neighbourhood of $\Sigma \cup \alpha$. Then Σ' is again a sphere, it is separating, and one of the capped components of $M \setminus \Sigma'$ is $S^2 \times S^1$ or $S^2 \tilde{\times} S^1$, so M has one such connected summand. The very basis of this argument breaks down for orbifolds, because if Σ is a spherical 2-orbifold of cyclic or vertex type, then Σ' is not spherical at all (see Fig. 3). We will devote Section 5 to showing that indeed there is a wealth of 3-orbifolds which contain non-separating spherical 2-orbifolds but cannot be expressed as a non-trivial connected sum.

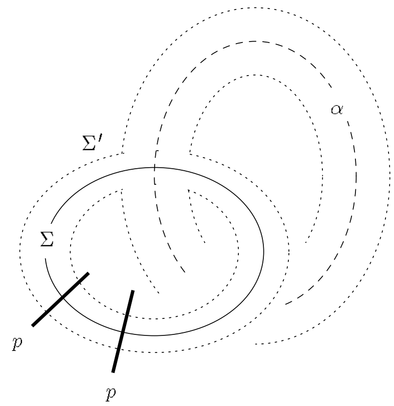

Figure 3: A non-separating spherical 2-orbifold may not determine a separating one.

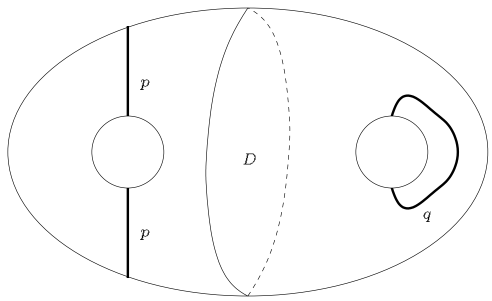

Figure 4: Compression along an ordinary disc does not preserve essentiality.

Turning to points (b) and (c) of the Introduction, recall that a family $\mathcal{S}$ of spherical 2-suborbifolds of a 3-orbifold X is called *essential* if no component of $X \setminus \mathcal{S}$ is a punctured discal 3-orbifold. To prove finiteness of a maximal essential family using the analogue of orbifolds of Haken's normal surface theory, one should be able to prove that the property of being essential is stable under the "normalization moves." Now one can see (as we will in Section 4), that the normalization moves boil down to isotopy and compression along ordinary discs (*i.e.* discs without singular points). And we have the following:

Example 2.1. Let a 3-orbifold X contain a system $\mathcal{S}$ of three cyclic spherical 2-suborbifolds or orders p, p, and q, for some $p, q \geqslant 2$. Suppose that one of the components of $X \setminus \mathcal{S}$ is as shown in Fig. 4 (an ordinary connected sum

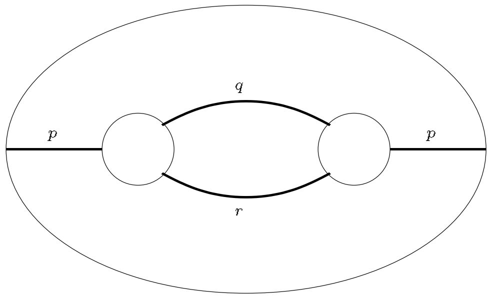

Figure 5: The capped components of the complement of a maximal essential system need not be irreducible.

of $\mathbb{S}^3_{\rm c}(p)$ and $\mathbb{S}^3_{\rm c}(q)$, with two cyclic punctures of order p and one of order q). Suppose that the other components of $X\setminus\mathcal{S}$ are once-punctured irreducible non-spherical 3-orbifolds. Let $\mathcal{S}'$ and $\mathcal{S}''$ be the systems obtained from $\mathcal{S}$ by compression along the ordinary disc D also shown in Fig. 4. *Then X does not contain any bad 2-suborbifold, every spherical 2-suborbifold of X is separating, $\mathcal{S}$ is essential, $\mathcal{S}'$ and $\mathcal{S}''$ are systems of spherical 2-orbifolds, and they are both inessential in X.*

The failure of the direct orbifold analogue of normal surfaces to establish existence of a maximal essential system may appear to be a technical difficulty, only calling for a different proof of existence. But the next example, which refers to point (c) in the Introduction, shows that the notion itself of *essential* system does not do the job it was designed for:

Example 2.2. Let a 3-orbifold X contain a system $\mathcal{S}$ of one cyclic and two vertex spherical 2-suborbifolds or orders p and p,q,r (twice) for some admissible $p,q,r\geqslant 2$. Suppose that one of the components of $X\setminus\mathcal{S}$ is as shown in Fig. 5, that is, $\mathbb{S}^3_{\rm v}(p,q,r)$ with two vertex and one order-p cyclic punctures). Let $\mathcal{S}'$ be given by the two vertex components of $\mathcal{S}$. Suppose that the other components of $X\setminus\mathcal{S}$ are once-punctured irreducible non-spherical 3-orbifolds. Then X *does not contain any bad 2-suborbifold, every spherical 2-suborbifold of X is separating, $\mathcal{S}'$ is a maximal essential system spherical 2-orbifolds in X, but not all the capped components of $X\setminus\mathcal{S}'$ are irreducible.* To understand the last two assertions, note that $\mathcal{S}$ is not essential, because the 3-orbifold in Fig. 5 is indeed punctured discal, but the cyclic component of $\mathcal{S}$ is essential after cutting along $\mathcal{S}'$ and capping, because it corresponds to a *cyclic* connected sum with $\mathbb{S}^3_{\rm v}(p,q,r)$, which is non-trivial.

It is perhaps worth pointing out the basic fact underlying the previous example, namely that *a punctured spherical* 3*-orbifold is a discal one if and only if the puncture has the same type as the original* 3*-orbifold.*

Concerning uniqueness (point (d) of the Introduction), all we had to say about its statement was said in Section 1. About its proof, we only note that Milnor's one for manifolds was based on the same method proving existence, *i.e.* on the notion of essential spherical system, so of course our proof will be somewhat different.

3 Existence of efficient splitting systems and algorithmic aspects

In this section we prove the existence parts of Theorems 1.9 and 1.7. We do so by showing that the process described in Theorem 0.2 can be carried out with all the components of the splitting system simultaneously normal with respect to a fixed triangulation.

Triangulations and normal 2-orbifolds We slightly relax the definition given in [2], and define a *triangulation* of a 3-orbifold X as a triangulation of $|X|$ which contains $S(X)$ as a subcomplex of its one-skeleton. We then define a 2-suborbifold of X to be normal with respect to a triangulation if it intersects each tetrahedron in a union of squares and triangles, just as in the case of a surface in a triangulated manifold. The following fundamental result holds:

Proposition 3.1. *Let X be a triangulated* 3*-orbifold. Then there exists a number m such that, if $\mathcal{F}$ is a normal surface with respect to the given triangulation and $\mathcal{F}$ has more than m connected components, then at least one of the components of $X \setminus \mathcal{F}$ is a product orbifold.*

This result is established exactly as in the case of manifolds, so we omit the proof. One only needs to notice that if a region of $X \setminus \mathcal{F}$ is a product as a manifold then it is also a product as an orbifold.

Normalization moves We now recall that, in a manifold with a fixed triangulation T, there is a general strategy to replace a surface by a normal one, as explained for instance in [7]. When the surface is a sphere Σ, this strategy consists (besides isotopy relative to T) of three "normalization" moves. *The first (respectively, second) move consists in compressing Σ along a disc D contained in the interior of a tetrahedron (respectively, triangle) of*

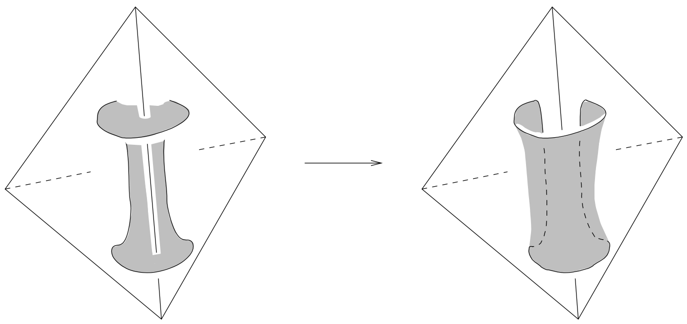

Figure 6: Third normalization move.

T, thus getting two spheres Σ' and Σ'', and choosing either Σ' or Σ''. The third move is used to remove the double intersections of the normal discs with the edges, and it is illustrated in Fig. 6. As a matter of fact, this move can also be seen as compression along a disc followed by the choice of one of the resulting spheres. More precisely, one compresses along the disc in Fig. 6-right and one discards the sphere given by the union of the two discs in Fig. 6. Therefore, *the third move consists of the compression along a disc disjoint from the 1-skeleton of T, followed by the elimination of a sphere which is the boundary of a regular neighbourhood of an interior point of an edge of T.*

Stepwise normalization As already mentioned, to prove the existence of an efficient splitting system we will construct a system of spherical 2-orbifolds along the steps of Theorem 0.2, making sure each new component of the system is normal with respect to a fixed triangulation of the original orbifold. To show that this is possible, we start with an easy lemma, that we will use repeatedly below. For its statement, we stipulate that a spherical 2-orbifold of cyclic type is *more complicated* than one of ordinary type, and a that a spherical 2-orbifold of vertex type is *more complicated* than one of ordinary or cyclic type.

Lemma 3.2. *Let Δ be a punctured discal 3-orbifold, and let Σ be a component of $\partial\Delta$. Suppose that no component of $\partial\Delta$ is more complicated than Σ. Then capping all the components of $\partial\Delta$ except Σ we get a discal 3-orbifold.*

The next two results are the core of our argument on stepwise normalization. The first one will be used again in the next section. From now on we fix a 3-orbifold X and we suppose that X does not contain any bad 2-suborbifold, and that every spherical 2-suborbifold of X is separating.

Proposition 3.3. *Let $\mathcal{S}$ be an essential system of spherical 2-suborbifolds of X. Let Y be a component of $X \setminus \mathcal{S}$ and Σ be a spherical 2-suborbifold of Y. Suppose that:*

- *Σ is essential in $\widehat{Y}$.*
- *All the components of $\mathcal{S}$ are at most as complicated as Σ.*

Then $\mathcal{S} \cup \Sigma$ is essential.

Proof. Suppose that some component Z of $X \setminus (\Sigma \cup \mathcal{S})$ is punctured discal. The assumptions that $\mathcal{S}$ is essential and that Σ is separating imply that Z is contained in Y and Σ is a boundary component of Z. By Lemma 3.2 and the second assumption, capping all the components of ∂Z but Σ we get a discal 3-orbifold, which implies that Σ bounds a discal 3-orbifold in $\widehat{Y}$: a contradiction to the first assumption. $\square$

Proposition 3.4. *Fix a triangulation T of X. Let $\mathcal{S}$ be an essential system of spherical 2-suborbifolds in normal position with respect to T. Let Y be a component of $X \setminus \mathcal{S}$ and Σ be a spherical 2-suborbifold of Y. Suppose that:*

- *Σ is essential in $\widehat{Y}$.*
- *All the spherical 2-suborbifolds of $\widehat{Y}$ strictly less complicated than Σ are inessential in $\widehat{Y}$.*

Then Σ can be replaced by a spherical 2-orbifold satisfying the same properties and such that $\mathcal{S} \cup \Sigma$ is in normal position with respect to T.

Proof. We must show that a normalization move applied to Σ can be realized without touching $\mathcal{S}$, without changing the type of Σ, and preserving the property that Σ is essential $\widehat{Y}$. The first assertion is a straight-forward consequence of the description of the normalization moves given above. For the second and third assertions, we first note that in all cases the compression of a normalization move takes place along a disc disjoint from the 1-skeleton of T, whence along an ordinary (non-singular) disc D which meets $\Sigma \cup \mathcal{S}$ in $\partial D \subset \Sigma$ only. To fix the notation, let D' and D'' be the discs bounded by ∂D on Σ, and define $\Sigma' = D \cup D'$ and $\Sigma'' = D \cup D''$. By the assumptions that X contains no bad 2-suborbifolds and that D is an ordinary disc, we can (and will) assume up to permutation that Σ'' is an ordinary sphere, We now claim that Σ' or Σ'' is essential in $\widehat{Y}$. Before proving the claim, let us see how it implies the conclusion. Since Σ'' is ordinary, we have $\Sigma' \cong \Sigma$. If Σ' is essential, we are done. Otherwise Σ'' is essential (and ordinary), so Σ is also ordinary by the second assumption, and we can switch Σ' and Σ''.

To prove the claim, suppose by contradiction that Σ' and Σ'' are both inessential in $\widehat{Y}$. Since Σ'' is ordinary, we deduce that Σ bounds in $\widehat{Y}$ either a discal 3-orbifold glued along a 2-disc to an ordinary 3-disc, or a 3-orbifold obtained from a discal one by removing an ordinary 3-disc which intersects the boundary in a 2-disc. The result is in both cases a discal 3-orbifold (isomorphic to the original one), whence the conclusion. □

The result just proved easily implies the following:

Corollary 3.5. *If X is triangulated then one can carry out the process of Theorem 0.2 with the splitting system $\mathcal{S}$ in normal position at each step.*

This corollary, together with Propositions 3.1 and 3.3 immediately yields the result we were heading for:

Corollary 3.6. *X admits finite efficient splitting systems.*

As promised in the Introduction, we explain here why it is crucial for us to assume that all the spherical 2-suborbifolds of X are separating. First of all, the definition itself of essential system would not be appropriate without this assumption, because a non-separating spherical 2-orbifold could constitute an inessential system. In addition, the assumption was used in a fundamental way in the proof of Proposition 3.3.

Proposition 3.4 also has the next consequence, which proves the last assertion of Theorem 0.2:

Corollary 3.7. *Let Z be a triangulated orbifold and let Σ be an essential spherical 2-suborbifold of Z. Suppose that all the spherical 2-suborbifolds of Z strictly less complicated than Σ are inessential. Then there exists an essential spherical 2-suborbifold of Z of the same type as Σ and in normal position with respect to the fixed triangulation.*

4 Uniqueness of the splitting

In this section we conclude the proofs of Theorems 1.9, 1.7, and 0.2. As above we fix a 3-orbifold X and we suppose that X does not contain any bad 2-suborbifold, and that every spherical 2-suborbifold of X is separating.

Gluing of discal 3-orbifolds Our uniqueness result uses the following two easy technical lemmas. We include a proof of the second one, leaving to the reader the even easier proof of the first one.

Lemma 4.1. *Let W be a 3-orbifold and let $D \subset \partial|W|$ be a disc. Let Δ be a punctured discal 3-orbifold. Let Σ be a component of $\partial\Delta$ of the same type as Δ. Split $|\Sigma|$ along a loop into two discs D' and D'', and suppose that D, D', D'' are isomorphic to each other, with at most one singular point. Let W' be obtained from W by gluing Δ along a homeomorphism $D' \to D$. Then:*

1. *W' is punctured discal if and only if W is.*
2. *$\widehat{W'} = \widehat{W}$.*

Lemma 4.2. *Let W be a 3-orbifold with a spherical boundary component Σ. Let Δ be a punctured discal 3-orbifold with at least two boundary components Σ', Σ'' isomorphic to Σ, and all other components not more complicated than Σ' and Σ''. Let W' be obtained from W by gluing Δ along a homeomorphism $\Sigma' \to \Sigma$. Then:*

1. *W' is punctured discal if and only if W is.*
2. *$\widehat{W'} = \widehat{W}$.*

Proof. We begin by (2). Since Σ' and Σ'' are at least as complicated as all the other components of $\partial\Delta$, capping the components different from Σ' and Σ'' we get a product orbifold, which implies the conclusion at once.

Let us turn to (1). To be discal for a 3-orbifold Z means that $\widehat{Z}$ is some $\mathbb{S}^3_{\mathrm{t}}$ and some component of ∂Z, or more precisely the most complicated one, is $\mathbb{S}^2_{\mathrm{t}}$. Now we already know that $\widehat{W'} = \widehat{W}$, and the assumptions imply that the most complicated component of $\partial W'$ has the same type as the most complicated component of ∂W, whence the conclusion. $\square$

Finiteness and uniqueness: scheme of the proof To conclude the proofs of the results stated so far we must establish the finiteness of the splitting process of Theorem 0.2 and the uniqueness part of Theorems 0.2, 1.7, and 1.9. The former easily follows from the next result, to be proved below, and Proposition 3.3.

Proposition 4.3. *Let $\mathcal{S}$ and $\mathcal{S}'$ be finite systems of spherical 2-suborbifolds of X. Denote by n_{t} (respectively, n'_{t}) the number of components of $\mathcal{S}$ (respectively, $\mathcal{S}'$) of type* t. *Suppose that $\mathcal{S}$ is coirreducible and $\mathcal{S}'$ is essential. Then $(n'_{\mathrm{o}}, n'_{\mathrm{c}}, n'_{\mathrm{v}}) \leqslant (n_{\mathrm{o}}, n_{\mathrm{c}}, n_{\mathrm{v}})$ in lexicographical order.*

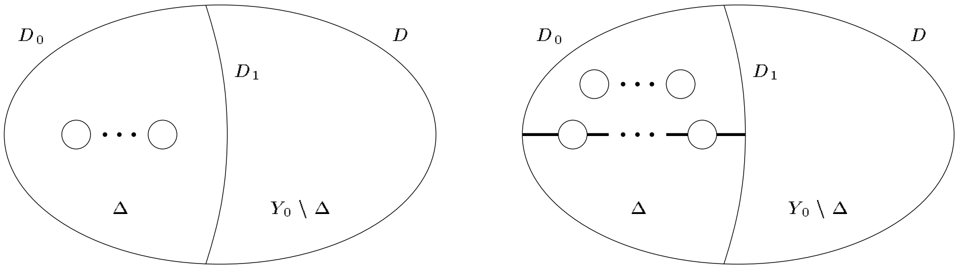

Figure 7: The two instances where move α applies, with $\widehat{\Delta} = \mathbb{S}^3_{\mathrm{o}}$ (left) and $\widehat{\Delta} = \mathbb{S}^3_{\mathrm{c}}$ (right).

Uniqueness easily follows from the next result (for Theorem 0.2 one only has to note that the final $\mathcal{S}$ is obviously coirreducible, and essential by Proposition 3.3):

Proposition 4.4. *Let $\mathcal{S}$ and $\mathcal{S}'$ be efficient spherical splitting systems for X. Then $\mathcal{S} \cong \mathcal{S}'$ and $(X \setminus \mathcal{S})\widehat{\ } \cong (X \setminus \mathcal{S}')\widehat{\ }$ as abstract collections of* 2- *and* 3-*orbifolds respectively.*

Moves for spherical systems We will prove Propositions 4.3 and 4.4 in a unified manner. The underlying idea in both cases is that by certain moves we can modify $\mathcal{S}$ into a superset of $\mathcal{S}'$, leaving unaffected the properties of $\mathcal{S}$ and the topological types of $\mathcal{S}$ and $(X \setminus \mathcal{S})\widehat{\ }$ as abstract orbifolds. There are two types of move we will need, that we now describe. These moves apply to any spherical system $\mathcal{S} \subset X$, under the appropriate assumptions.

Move α. Suppose we have a component Σ_0 of $\mathcal{S}$ and a disc D_1 properly embedded in a component Y_0 of $X \setminus \mathcal{S}$, with $\partial D_1 \subset \Sigma_0$. Let ∂D_1 split Σ_0 into discs D_0 and D, and suppose that D_0 and D_1 both have at most one singular point. Assume that the component containing D_0 of $Y_0 \setminus D_1$ is a punctured disc Δ, and that $\partial\Delta$ does not have components more complicated than $D_0 \cup D_1$. Then we replace $\Sigma_0 = D \cup D_0$ by $\Sigma_1 = D \cup D_1$. See Fig. 7.

Move β. Suppose we have a component Σ_0 of $\mathcal{S}$ and a spherical 2-orbifold Σ_1 contained in a component Y_0 of $X \setminus \mathcal{S}$, with $\Sigma_0 \subset \partial Y_0$. Suppose that the component containing Σ_0 of $Y_0 \setminus \Sigma_1$ is a punctured disc Δ, that Σ_0 and Σ_1 have the same type, and that no other component of $\partial\Delta$ is more complicated than Σ_0 and Σ_1. Then we replace Σ_0 by Σ_1. See Fig. 8.

Lemma 4.5. *Under any move of type α or β the following happens:*

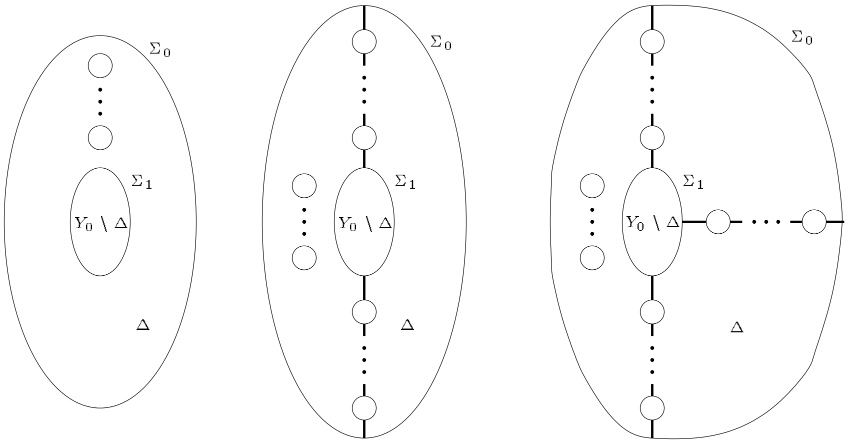

Figure 8: The three instances where move β applies, with $\widehat{\Delta} = \mathbb{S}^3_{\rm o}$ (left), $\widehat{\Delta} = \mathbb{S}^3_{\rm c}$ (centre), and $\widehat{\Delta} = \mathbb{S}^3_{\rm v}$ (right).

- *The type of $\mathcal{S}$ (including orders) does not change.*
- *The type of $(X \setminus \mathcal{S})\widehat{\ }$ does not change (in particular, coirreducibility is preserved).*
- *Essentiality is preserved.*

Proof. The first assertion is proved by direct examination of Figg. 7 and 8. To prove the second and third assertions we note the only components of $X \setminus \mathcal{S}$ modified by the move are Y_0 and the other component Z_0 incident to Σ_0, and these components are replaced by the components Y_1 and Z_1 incident to Σ_1. The relations between them are described as follows:

$$\text{move } \alpha: \quad Y_0 = Y_1 \cup_{D_1} \Delta, \quad Z_1 = Z_0 \cup_{D_0} \Delta;$$
$$\text{move } \beta: \quad Y_0 = Y_1 \cup_{\Sigma_1} \Delta, \quad Z_1 = Z_0 \cup_{\Sigma_0} \Delta.$$

It is now easy to check that in all cases we are under the assumptions of Lemmas 4.1 and 4.2, whence the conclusion at once. $\square$

Finiteness and uniqueness: conclusion As already mentioned, to prove Propositions 4.3 and 4.4 we want to replace $\mathcal{S}$ by a superset of $\mathcal{S}'$. The first step is as follows:

Proposition 4.6. *Suppose that $\mathcal{S}$ is coirreducible and $\mathcal{S}'$ is essential. Then, up to applying moves α and isotopy to $\mathcal{S}$, we can suppose that $\mathcal{S} \cap \mathcal{S}' = \emptyset$.*

Proof. Let us first isotope $\mathcal{S}$ so that it is transversal to $\mathcal{S}'$. Of course it is sufficient to show that if $\mathcal{S} \cap \mathcal{S}'$ is non-empty then the number of its components can be reduced by a move α and isotopy.

Suppose some component Σ' of $\mathcal{S}'$ meets $\mathcal{S}$, and note that the finite set of loops $\mathcal{S} \cap \Sigma'$ bounds at least two innermost discs in Σ'. Since Σ' has at most three singular points, we can choose an innermost disc D_1 which contains at most one singular point. Let Y_0 be the component of $X \setminus \mathcal{S}$ which contains D_1, let Σ_0 be the component of $\mathcal{S}$ containing ∂D_1, and denote by D_0 and D the discs into which ∂D_1 cuts Σ_0, with D_0 at most as complicated as D.

The assumption that X contains no bad 2-suborbifolds easily implies that $D_0 \cup D_1$ is a spherical 2-orbifold of ordinary or cyclic type. Coirreducibility of $\mathcal{S}$ then implies that $D_0 \cup D_1$ bounds an ordinary or cyclic discal 3-orbifold B in $\widehat{Y}_0$. Now we either have $B \cap D = \emptyset$ or $B \supset D$. In the first case it is easy to see that we are precisely in the situation where we can apply a move α, after which a small isotopy allows to reduce the number of components of $\mathcal{S} \cap \mathcal{S}'$. In the second case $D_1 \cup D$ bounds in $\widehat{Y}_0$ one of the manifolds obtained from B by cutting along D. Since B is an ordinary or cyclic disc we deduce that D is isomorphic to D_1 and that $D \cup D_1$ bounds a ball in $\widehat{Y}_0$ disjoint from D_0. We can then interchange the roles of D and D_0 and apply α also in this case, whence the conclusion. □

Propositions 4.3 and 4.4 now easily follow from Proposition 4.6 and the next result:

Proposition 4.7. *Suppose that $\mathcal{S}$ is coirreducible and $\mathcal{S}'$ is essential. Assume that each component of $\mathcal{S}'$ is either contained in $\mathcal{S}$ or disjoint from $\mathcal{S}$. Then:*

- *Up to applying moves β and isotopy to $\mathcal{S}$, we can suppose that each ordinary component of $\mathcal{S}'$ is contained in $\mathcal{S}$.*
- *If $\mathcal{S}$ and $\mathcal{S}'$ have the same ordinary components, up to applying moves β and isotopy to $\mathcal{S}$, we can suppose that each cyclic component of $\mathcal{S}'$ is contained in $\mathcal{S}$.*
- *If $\mathcal{S}$ and $\mathcal{S}'$ have the same ordinary and cyclic components, up to applying moves β and isotopy to $\mathcal{S}$, we can suppose that each vertex component of $\mathcal{S}'$ is contained in $\mathcal{S}$.*

Proof. In all the cases, we suppose that a certain component Σ_1 of $\mathcal{S}'$ of the appropriate type is not contained $\mathcal{S}$, and we show that a move β can be applied to $\mathcal{S}$ which replaces some component Σ_0 by Σ_1. The conclusion then easily follows by iteration. In all cases we denote by Y_0 the component of $X \setminus \mathcal{S}$ which contains Σ_1, and we choose Σ_0 among the components of ∂Y_0.

We start with the first assertion, so we assume Σ_1 is ordinary. Coirreducibility of $\mathcal{S}$ implies that Σ_1 bounds an ordinary 3-disc B in $\widehat{Y}_0$. Let us choose a component of $\mathcal{S}'$ contained in B, not contained in $\mathcal{S}$, and innermost in B with respect to these properties, and let us rename it Σ_1. (Note that the new Σ_1 may not be parallel to the old one in X, even if it is in B). Of course Σ_1 is again an ordinary sphere, and it bounds in Y a punctured ordinary disc Δ such that all the components of $\partial\Delta$ except Σ_1 belong to $\mathcal{S}$. Essentiality of $\mathcal{S}'$ now implies that there must be a component Σ_0 of $\partial\Delta$ which does not belong to $\mathcal{S}'$. We have created a situation where a move β can be applied, so we are done.

The proof is basically the same for the second and third assertions. We only need to note that B always has the same type as Σ_1 and that the component Σ_0 of $\partial\Delta$ not belonging to $\mathcal{S}'$ cannot be of type strictly less complicated than Σ_1, by the assumption that $\mathcal{S}$ and $\mathcal{S}'$ have the same components of type strictly less complicated than that of Σ_1. $\square$

5 Prime non-irreducible 3-orbifolds

In this section we treat point (d) of the Introduction, showing that there are infinitely many 3-orbifolds with non-separating but no essential separating spherical 2-suborbifolds. Our examples arose from discussions with Luisa Paoluzzi.

As in the case of manifolds, one can define a prime orbifold as one that cannot be expressed as a non-trivial connected sum, *i.e* an orbifold in which every *separating* spherical 2-suborbifold bounds a discal 3-suborbifold. Of course every irreducible orbifold is also prime, but we will show now that there are infinitely many primes which are not irreducible. This appears as a sharp difference between the orbifold and the manifold case (in which the only non-irreducible primes are the two S^2-bundles over S^1). This difference represents a serious obstacle to promoting Theorem 1.7 to a unique decomposition theorem for orbifolds into prime ones.

Proposition 5.1. *Let X be a 3-orbifold bounded by a sphere with 4 cone points of the same order p. Assume that:*

- *X is irreducible.*
- *X is not the 3-ball with two parallel unknotted singular arcs.*

- *In X there is no sphere with three singular points which meets ∂X in a disc with two singular points.*

Consider the orbifold with support $S^2 \times S^1$ and singular set given by two circles of order p parallel to the factor S^1. Remove from it a regular neighbourhood of an arc which joins the singular components within a level sphere $S^2 \times \{*\}$, and call Y the result. Let Z be obtained by attaching X and Y along their boundary spheres, matching the cone points. Then Z is prime but not irreducible.

Proof. A level sphere $S^2 \times \{*\}$ contained in Y gives a non-separating spherical 2-suborbifold of Z, so Z is not irreducible.

Let us consider a separating spherical 2-orbifold Σ in Z, and the intersection of Σ with the sphere Σ_0 along which X and Y have been glued together. If $\Sigma \cap \Sigma_0$ is empty then Y bounds a spherical 3-orbifold, because X and Y are irreducible. Let us then assume $\Sigma \cap \Sigma_0$ to be transverse and minimal up to isotopy of Σ. Considering the pattern of circles $\Sigma \cap \Sigma_0$ on Σ we see that Σ contains at least two innermost discs, which are therefore contained either in X or in Y. Since Σ has at most three singular points, we can find one such disc D with $i \leqslant 1$ points. Now consider the loop ∂D on Σ_0, and recall that Σ_0 has 4 singular points. Then ∂D encircles $e \leqslant 2$ singular points. We will now show that both if $D \subset X$ and if $D \subset Y$ the conditions $i \leqslant 1$ and $e \leqslant 2$ are impossible, whence the conclusion.

Suppose first that $D \subset X$. Using the fact that X is irreducible, the cases $e = i = 0$ and $e = i = 1$ both contradict the minimality of $\Sigma \cap \Sigma_0$. The cases $e = 0,\ i = 1$ and $e = 1,\ i = 0$ both lead to a bad 2-suborbifold of X, which does not exist. In the case $e = 2,\ i = 0$, using irreducibility twice, we see that X must be the 3-ball with two parallel unknotted singular arcs, but this orbifold was specifically forbidden, so indeed we get a contradiction. The case $e = 2,\ i = 1$ was also specifically forbidden.

Suppose now that $D \subset Y$. Recalling that D is contained in a sphere which is separating in Z, we see that D cobounds a ball with a disc contained in ∂Y. Using this fact we can prove again that all the possibilities with $e \leqslant 2$ and $i \leqslant 1$ are impossible. As above, $e = i = 0$ and $e = i = 1$ contradict minimality. The case $e = 2,\ i = 0$ would imply that one of the singular arcs of Y can be homotoped to ∂Y, which is not the case. The other cases are absurd for similar reasons. $\square$

Examples Whenever $k \geqslant 1$ and the cone angles are chosen in an admissible fashion, the "stairway" of Fig. 9 contained in the 3-ball provides an

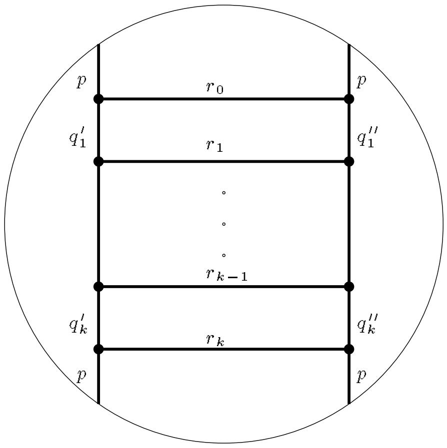

Figure 9: An example or orbifold X such that $Z = X \cup Y$ is prime but not irreducible.

example of orbifold X satisfying the assumptions of Proposition 5.1. Note that for such an X the corresponding prime non-irreducible orbifold Z is always supported by $S^2 \times S^1$, so $|Z|$ is prime as a manifold. But we can also construct examples of X to which Proposition 5.1 applies, leading to an orbifold Z with non-prime $|Z|$. For instance, one can take an orbifold constructed exactly as Y was constructed from $S^2 \times S^1$, starting instead with $F \times S^1$, where F is a closed orientable surface F of positive genus.

6 3-orbifolds with ordinary and cyclic splitting

Our interest in the spherical splitting of 3-orbifolds arose from our program of extending Matveev's complexity theory [5, 6] from the manifold to the orbifold context. We actually succeeded [9] in generalizing the definition and some of the most significant results, investigating in particular the behaviour under connected sum. However it turns out that one does not have in general plain additivity as in the manifold case. We can only prove additivity when there are no vertex connected sums, and we have that complexity is only additive up to a certain correction summand on cyclic connected sums involving at least one purely cyclic singular component. For this reason in this section we prove that the number of such sums is well-defined:

Proposition 6.1. *Suppose that a 3-orbifold X is a connected sum of irreducible orbifolds without vertex connected sums, and fix an efficient realization ρ of X as $X_0\#\ldots\#X_n$. For all $p \geqslant 2$ let $\nu_\rho(p)$ be the number* *of p-cyclic connected sums in ρ involving at least one singular component without vertices. Then $\nu_\rho(p)$ is independent of ρ, so it is a function of X only.*

Proof. Two different realizations of X are related by the following operations:

- Reordering of the X_i's, without modification of the efficient system of spherical 2-orbifolds along which X is split.
- Modification of the splitting system according to one of the moves α and β described in Section 4 (see Figg. 7 and 8).

We will show that $\nu_\rho(p)$ remains unchanged under both operations.

To deal with the first operation, we define a modified version T_p of the tree associated to the realization $X = X_0\#\ldots\#X_n$. To construct T_p we associate to each X_i one node for each cyclic component of $S(X_i)$ and one for the union of the non-cyclic components of $S(X_i)$. We denote by C the set of nodes of cyclic type and by N the set of nodes of non-cyclic type. We now define T_p by taking an edge for each p-cyclic connected sum, with the ends of the edge joining the singular components involved in the sum.

The ideas behind the construction of T_p are as follows. First, we cannot create a new p-cyclic singular component by performing an ordinary connected sum or a q-cyclic connected sum for $q \neq p$. Second, a cyclic connected sum between two cyclic components gives one cyclic component, while a cyclic connected sum between a non-cyclic component and any other one gives one or two non-cyclic components. With these facts in mind, it is very easy to see that the invariance of $\nu_\rho(p)$ under the first move is implied by the following graph-theoretical statement: *Let T be a tree with set of nodes $C \sqcup N$. Let $\sigma = (e_1, \ldots, e_n)$ be an ordering of the edges of T. For $k = 1, \ldots, n$ define $\alpha(k)$ to be 0 if both the ends of e_k can be joined to nodes in N by edges in $e_1, \ldots, e_{k-1}$ only, and to be 1 otherwise. Then $\sum_{k=1}^{n} \alpha(k)$ is independent of σ.*

To prove this fact, we first reduce to an easier statement. We begin by noting that the vertices in N having valence more than 1 can be blown up as in Fig. 10 without affecting the sum in question. This allows us to assume that all the nodes in N are "external," *i.e.* 1-valent. Moreover, using induction, we can suppose that N consists precisely of the external nodes, because if an edge e has an end which is external and belongs to C then $\alpha(e) = 1$ whatever the position of e in the ordering, and the value of α on the remaining edges does not depend on the position of e in the ordering.

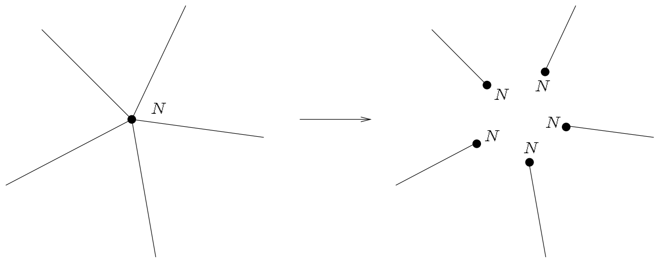

Figure 10: This move does not change the sum of the values of α.

Switching from $\alpha(e)$ to $\beta(e) = 1-\alpha(e)$ it is then sufficient to establish the following result. The proof we present here is due to Gwénaël Massuyeau.

Lemma 6.2. *Let T be a tree. Let E denote the set of external nodes of T. Let $\sigma = (e_1, \ldots, e_n)$ be any ordering of the edges of T. Define $\beta(k)$ to be 1 if the ends of e_k can be joined to E through edges in $e_1, \ldots, e_{k-1}$ only, and to be 0 otherwise. Then $\sum_{k=1}^{n} \beta(k)$ is independent of σ, and precisely equal to $\#E - 1$.*

Proof. Attach a circle to T by gluing the nodes in E to $\#E$ arbitrarily chosen points on the circle, and denote by Y the result. Since T is a tree, we have $\chi(Y) = 1 - \#E$. Let $\sigma = (e_1, \ldots, e_n)$ be an ordering of the edges of T, let Y_k be the union of the circle with $e_1, \ldots, e_k$, and let $\widetilde{Y}_k$ be the connected component of Y_k which contains the circle. Of course $\widetilde{Y}_0$ is the circle and $\widetilde{Y}_k = Y$. We now have the identity

$$\chi(\widetilde{Y}_k) = \chi(\widetilde{Y}_{k-1}) - \beta(k)$$

which is easily established by considering separately the cases where e_k has 0, 1, or 2 ends on $\widetilde{Y}_{k-1}$, and using again the fact that T is a tree. The identity implies that

$$\sum_{k=1}^{n} \beta(k) = \chi(\widetilde{Y}_0) - \chi(\widetilde{Y}_k) = \#E - 1.$$

□

To conclude the proof of Proposition 6.1 we must now show that $\nu_\rho(p)$ is not affected by the moves α and β. Knowing the independence of the ordering, this is actually very easy. For both the moves, we have two realizations ρ_0 and ρ_1 of X leading to splitting systems $\mathcal{S}_0$ and $\mathcal{S}_1$ which differ for one component only, so $\mathcal{S}_0 = \{\Sigma_0\} \cup \mathcal{S}$ and $\mathcal{S}_1 = \{\Sigma_1\} \cup \mathcal{S}$. We then compute $\nu_{\rho_j}(p)$ by performing first the connected sum which corresponds to Σ_j. From Figg. 7 and 8 one sees that the connected sums along Σ_0 and Σ_1 give the same contributions to ν. After performing them, all other sums in ρ_0 and ρ_1 are the same, which eventually proves the proposition. □

Remark 6.3. Proposition 6.1 is false if one allows also vertex connected sums in the realization of X, because a vertex connected sum can create singular components without vertices.

References

[1] M. Boileau – B. Leeb – J. Porti, *Geometrization of 3-dimensional orbifolds*, `math.GT/0010184`.

[2] M. Boileau – S. Maillot – J. Porti, "Three-dimensional orbifolds and their geometric structures," Panoramas et Synthèses Vol. 15, Société Mathématique de France, Paris, 2003.

[3] J. Hempel, "3-Manifolds," Ann. of Math. Studies Vol. 86, Princeton University Press, Princeton, N. J., 1976.

[4] D. Cooper – C. Hodgson – S. Kerkhoff, "Three-dimensional orbifolds and cone-manifolds," MSJ Memoirs Vol. 5, Mathematical Society of Japan, Tokyo, 2000.

[5] B. Martelli, *Complexity of 3-manifolds*, `math.GT/0405250`.

[6] S. V. Matveev, *Complexity theory of three-dimensional manifolds*, Acta Appl. Math. **19** (1990), 101-130.

[7] S. V. Matveev, "Algorithmic topology and classification of 3-manifolds," Algorithms and Computation in Mathematics Vol. 9. Springer-Verlag, Berlin, 2003.

[8] G. Perelman, *Ricci flow with surgery on three-manifolds*, `math.DG/0303109`.

[9] C. Petronio, *Complexity of 3-orbifolds*, in preparation.

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

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