Complexity of 3-orbifolds

Carlo Petronio*

October 21, 2004

Abstract

We extend Matveev's theory of complexity for 3-manifolds, based on simple spines, to (closed, orientable, locally orientable) 3-orbifolds. We prove naturality and finiteness for irreducible 3-orbifolds, and, with certain restrictions and subtleties, additivity under orbifold connected sum. We also develop the theory of handle decompositions for 3-orbifolds and the corresponding theory of normal 2-suborbifolds.

MSC (2000): 57M99.

Introduction

The aim of this paper is to extend Matveev's theory of complexity from 3-manifolds to 3-orbifolds, and in particular, after giving the appropriate definition, to prove the following properties, established for manifolds in [4]:

- Naturality and finiteness for irreducible orbifolds;
- Additivity under orbifold connected sum.

To quickly summarize our results, we anticipate that we have fully achieved the former of these tasks, with the most natural definition of complexity. Concerning the latter task, we have shown additivity to hold in a strict (and subtle) fashion for two of the three possible types of connected sum, and we have established a two-sided linear estimate for the third type.

We now give a more detailed account of our results. We always consider closed, orientable, locally orientable 3-orbifolds. To define the *complexity* $c(X)$ of such an X we consider the set of simple polyhedra P embedded in

*Research supported by the INTAS Project "CalcoMet-GT" 03-51-3663

$|X|$ so that P is transversal to the singular set $S(X)$ and $X \setminus P$ consists of disjoint open discal 3-orbifolds. Then we take the minimum over all P of the number of vertices of P plus a contribution for each intersection point between P and $S(X)$, the contribution being $p-1$ if the order of $S(X)$ at the point is p. We can then state the naturality and finiteness properties proved below:

Theorem 0.1. *With certain well-understood exceptions, if P realizes the complexity of an irreducible orbifold X then P is a special polyhedron, and dual to P there is a triangulation of X. The number of exceptional orbifolds of each given complexity is finite.*

Corollary 0.2. *For each positive integer n the set of irreducible orbifolds X such that $c(X) \leqslant n$ is finite.*

Turning to additivity, first recall that the operation of connected sum of two manifolds or orbifolds consists in making a puncture in each of the given objects and then gluing the resulting boundary spheres. For manifolds, there is only one way to make the puncture, so the operation of connected sum is uniquely defined (at least in a connected and oriented context). In the analogue operation for orbifolds, one has to distinguish according to the nature of the point at which the puncture is made, which can be non-singular, singular but not a vertex of the singular locus, or a vertex of the singular locus. Depending on this nature, the operation of connected sum is called of *ordinary*, *cyclic*, or *vertex* type. Each type of operation has a corresponding identity element, given by the so-called *ordinary*, *cyclic*, and *vertex spherical* 3*-orbifold*. A connected sum of a certain type with the spherical 3-orbifold of the same type will be called *trivial*. The following result, proved in [6], is necessary to state our results on additivity.

Theorem 0.3. *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 connected sum of irreducible 3-orbifolds in such a way that, even after reordering the sums, there is no trivial sum. Any two such realizations involve the same irreducible summands and the same types of sums (including orders). Moreover, if the realization of X does not involve vertex connected sums, for all $p \geqslant 2$ the number $\nu(p)$ of p-cyclic connected sums involving at least one singular component without vertices is independent of the realization, so it is a function of X only.*

A realization of X as described in the previous statement will be called an *efficient connected sum*. The condition that there should not be trivial sums *even after reordering* is due to the existence of phenomena of the following type. If we first take the ordinary connected sum of some non-singular X with a cyclic spherical 3-orbifold, and then we take the cyclic connected sum of the result with some other Y, then both sums, taken in this order, are non-trivial, but the result is just the ordinary connected sum of X and Y, so the two-step sum is obviously inefficient. We now have our main statement about additivity:

Theorem 0.4. *Let a 3-orbifold X as in Theorem 0.3 be the efficient connected sum of irreducible orbifolds $X_1, \ldots, X_n$ without vertex connected sums. For $p \geqslant 2$ let $\nu(p)$ be as in Theorem 0.3. Then*

$$c(X) = \sum_{j=1}^{n-1} c(X_i) - \sum_{p=2}^{\infty} \nu(p) \cdot (p-1).$$

When the realization of X includes vertex connected sums we only have two-sided linear estimates on the complexity of X in terms of the complexity of its summands:

Theorem 0.5. *Let a 3-orbifold X as in Theorem 0.3 be the efficient connected sum of irreducible orbifolds $X_1, \ldots, X_n$. Then*

$$\frac{1}{4^{n-1}} \cdot \big(c(X_1) + \ldots + c(X_n)\big) \leqslant c(X) \leqslant 6^{n-1} \cdot \big(c(X_1) + \ldots + c(X_n)\big).$$

We mention that to establish the additivity properties of complexity we have developed an apparently new portion of the general theory of orbifolds, namely the notion of handle decomposition. The reader will find the details on this notion in Section 4, together with the analogue in the orbifold setting of Haken's theory of normal surfaces with respect to handle decompositions.

As a motivation for our interest in 3-orbifolds, we would like to recall here that orbifolds play a central rôle in 3-dimensional geometric topology, and in particular in Thurston's geometrization program, see [1].

1 Preliminaries and main definitions

In this section we briefly recall the notions of orbifold and of simple polyhedron, and we define the complexity of a 3-orbifold.

Local structure of orbifolds We will not cover here the general theory of orbifolds, referring the reader to the milestone [8], 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 and globally orientable.

Under these assumptions one can see that a 2-orbifold Σ consists of a compact orientable *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. In the sequel we will need the following easy result:

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* $p \neq q$*.*

We now introduce some notation and terminology repeatedly used below. We define $\mathbb{D}^3_{\mathrm{o}}$ to be D^3, the *ordinary discal* 3-orbifold, $\mathbb{D}^3_{\mathrm{c}}(p)$ 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 admissible orders p, q, r. We will call $\mathbb{D}^3_{\mathrm{c}}(p)$ and $\mathbb{D}^3_{\mathrm{v}}(p, q, r)$ respectively *cyclic discal* and *vertex discal* 3-orbifolds, and we will employ the shortened notation $\mathbb{D}^3_{\mathrm{c}}$ and $\mathbb{D}^3_{\mathrm{v}}$ to denote cyclic and vertex discal 3-orbifolds with generic orders.

We also define the *ordinary, cyclic*, and *vertex spherical* 2-orbifolds, 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

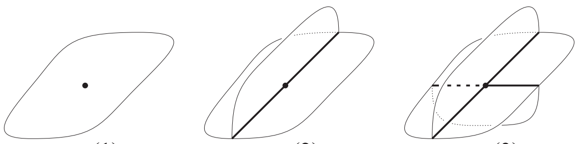

(1) (2) (3)

Figure 1: Local aspect of an almost-special polyhedron.

the corresponding discal 3-orbifolds $\mathbb{D}^3_{\rm o}$, $\mathbb{D}^3_{\rm c}(p)$, and $\mathbb{D}^3_{\rm v}(p,q,r)$. Finally, we introduce the *ordinary, cyclic*, and *vertex spherical* 3-orbifolds, denoted respectively by $\mathbb{S}^3_{\rm o}$, $\mathbb{S}^3_{\rm c}(p)$, and $\mathbb{S}^3_{\rm v}(p,q,r)$, as the 3-orbifolds obtained by mirroring the corresponding discal 3-orbifolds $\mathbb{D}^3_{\rm o}$, $\mathbb{D}^3_{\rm c}(p)$, and $\mathbb{D}^3_{\rm v}(p,q,r)$ in their boundary. The spherical 2- and 3-orbifolds with generic orders will be denoted by $\mathbb{S}^2_*$ and $\mathbb{S}^3_*$.

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 X is called *irreducible* if it does not contain any bad 2-suborbifold and every spherical 2-suborbifold of X is inessential (in particular, it is separating).

Simple and special polyhedra From now on we will employ the piecewise linear viewpoint, which is equivalent to the smooth one in dimensions 2 and 3. We will use the customary notions of PL topology without recalling them, see [7]. A *simple polyhedron* is a compact polyhedron P such that the link of each point of P can be embedded in the space given by a circle with three radii. In particular, P has dimension at most 2. Finite graphs and closed surfaces are examples of simple polyhedra. A point of a simple polyhedron is called a *vertex* if its link is precisely given by a circle with three radii. A regular neighbourhood of a vertex is shown in Fig. 1-(3). From the figure one sees that the vertices are isolated, whence finite in number. Graphs and surfaces do not contain vertices. The complexity $c(P)$ of P is the number of vertices that P contains.

Two more restrictive types of polyhedra will be used below. A simple polyhedron P is called *almost-special* if the link of each point of P is given

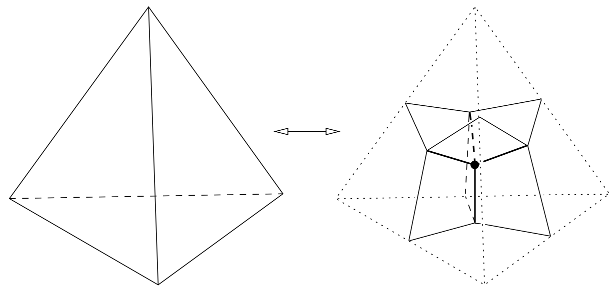

Figure 2: Duality between triangulations and special polyhedra.

by a circle with either zero, or two, or three radii. The local aspects of P are correspondingly shown in Fig. 1. The points of type (2) or (3) are called *singular*, and the set of singular points of P is denoted by $S(P)$. We will say that P is *special* if it is almost-special, $S(P)$ contains no circle component, and $P \setminus S(P)$ consists of open 2-discs.

Duality The following duality result, that we will state in the closed context only, is well-known. We call *triangulation* of a 3-manifold M a realization of M as a gluing of a finite number of tetrahedra along a complete system of simplicial pairings of the lateral faces. Note that we allow multiple and self-adjacencies of the tetrahedra, thus relaxing the traditional requirements for a triangulation used in PL topology.

Proposition 1.2. *The set of triangulations of a* 3*-manifold* M *corresponds bijectively to the set of special polyhedra* P *embedded in* M *in such a way that* $M \setminus P$ *is a union of open* 3*-discs. Both triangulations and special polyhedra are viewed up to isotopy, and the polyhedron corresponding to a triangulation is the* 2*-skeleton of the dual cellularization, as shown in Fig. 2.*

Spines and complexity Let X be a (closed, orientable, locally orientable) 3-orbifold. Let P be a simple polyhedron contained in $|X|$. We will say that P is a *spine* of X if the following holds:

- The intersections between P and $S(X)$ occur only at surface points of P and non-vertex points of $S(X)$, and they are transverse;
- If $U(P)$ is an open regular neighbourhood of P in $|X|$, each component of $X \setminus U(P)$ is isomorphic to one of the discal 3-orbifolds $\mathbb{D}^3_*$.

In the case of a manifold, *i.e.* for $S(X)=\emptyset$, we recover Matveev's condition that $X\setminus P$ consists of open 3-discs, *i.e.* that X minus some points collapses onto P. Note that, as opposed to the manifold case, the existence of a non-empty singular locus in an orbifold forces a spine to have some 2-dimensional stratum. From now on we will identify $X\setminus U(P)$ to $X\setminus P$, making a distinction only when it really matters. The reader will easily check that this choice does not cause any ambiguity.

If P is a spine of X as above, we define the following function, which depends not only on P as an abstract polyhedron, but also on its embedding in $|X|$ relative to $S(X)$:

$$c(P,S(X))=c(P)+\sum\left\{p-1:\ x\in P\cap S(X),\ \text{the order of } S(X) \text{ at } x \text{ is } p\right\}.$$

We now define $c(X)$, the *complexity* of X, as the minimum of $c(P,S(X))$ over all spines P of X. We note that c is always well-defined because every orbifold has simple spines: see for instance Section 2, where the relation is discussed between the special spines of a 3-orbifold X and the (suitably defined) triangulations of X.

We will say that a spine P of X is *minimal* if $c(P,S(X))=c(X)$ and every proper subpolyhedron of P which is also a spine of X is actually homeomorphic to P.

Remark 1.3. Each of the following may appear to be a more natural condition to include in the definition that P is *minimal*, besides the requirement that $c(P,S(X))=c(X)$:

(A) P does not collapse onto any proper subpolyhedron which is also a spine of X;

(B) P does not contain any proper subpolyhedron which is also a spine of X.

Condition (A) is however not satisfactory, because a minimal spine would remain minimal after a "bubble move" (the gluing of a 2-disc along the boundary of a 2-disc contained in P and disjoint from $S(X)$ and $S(P)$). This phenomenon already appears in the manifold case, where Matveev uses condition (B). However (B) indeed is not the right condition for orbifolds, for otherwise some orbifolds, *e.g.* $\mathbb{S}^3_{\mathrm{c}}(p)$, would not have minimal spines — see the proof of Lemma 2.2. One drawback of our definition of minimality is that it does not immediately imply that minimal spines exist, and it will require some efforts to prove that this is actually the case.

2 Minimal spines of irreducible orbifolds

In this section we prove that, with some well-understood exceptions, an irreducible orbifold has special minimal spines. This will imply that for any given n there is only a finite number of irreducible 3-orbifolds having complexity n, a result which opens the way to computer enumeration of orbifolds in order of increasing complexity, analogous to that carried out for manifolds by Martelli, Matveev, and the author, see [3].

We also give a natural notion of triangulation for an orbifold, and we show that a special spine is typically dual to a triangulation, which implies that our notion of complexity is a very natural one.

Singular arcs meet the spine once We begin with an easy result that we will use repeatedly both in this section and in Section 5.

Lemma 2.1. *Let P be a spine of an orbifold X such that $c(P, S(X)) = c(X)$. Let α be a connected component of $S(X)$ minus the vertices. Then $\alpha \cap P$ consists of precisely one point.*

Proof. At least one intersection point exists otherwise some component of $X \setminus P$ would contain either two vertices or a non-simply connected portion of $S(X)$. Suppose there are there are at least two intersection points, and let β be the open arc of α between two consecutive ones x_0 and x_1. Now β is contained in one of the open 3-discs B of which $X \setminus P$ consists. Let B' be the component of $X \setminus P$ which contains the portion of α past x_1. A priori we could have $B' = B$, but $S(B)$ consists of β only, which implies that $B' \neq B$. Therefore we can puncture P near x_1, getting a new spine P' of X. Moreover $c(P', S(X)) < c(P, S(X))$, a contradiction. □

Nuclear and minimal spines A simple polyhedron is called *nuclear* if it cannot be collapsed onto a proper subpolyhedron. Of course a minimal spine of a manifold is always nuclear. We will now show that orbifolds admit minimal spines, and that these spines are most often nuclear.

Lemma 2.2. *Let X be any orbifold and Y be the ordinary connected sum of X and $\mathbb{S}^3_{\rm c}(p)$. Then $c(Y) = c(X) + (p-1)$, and Y admits minimal spines if X does.*

Before proving the lemma we note that, if we apply it to $X = \mathbb{S}^3_{\rm o}$, we deduce that $c(\mathbb{S}^3_{\rm c}(p)) = p - 1$.

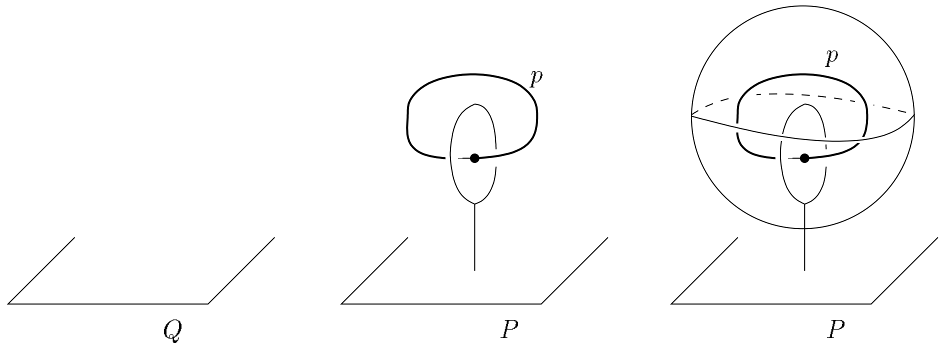

Figure 3: Spine of an orbifold after insertion of an extra unknotted singular component.

Proof. Suppose P is a spine of Y with $c(P, S(Y)) = c(Y)$. Let us denote by U_p the singular set of $\mathbb{S}^3_{\rm c}(p)$, an order-p unknot in the 3-sphere. Since $|Y| = |X|$ and $S(Y) = S(X) \sqcup U_p$ then P is also a spine of X, and

$$c(P, S(Y)) \geqslant c(P, S(X)) + (p-1),$$

because U_p must meet P in at least one point. This proves that $c(Y) \geqslant c(X)+(p-1)$. We will now show that, conversely, if Q is a spine of X, we can construct a spine P of Y such that $c(P, S(Y)) = c(Q, S(X)) + (p-1)$, which will imply the first assertion. We construct P so that it will be automatically minimal if Q is minimal, which will prove the second assertion. To do so, we distinguish according to whether there exists a component of $X \setminus P$ which is an ordinary 3-disc. If there is one, we construct P by attaching to Q a lollipop contained in this 3-disc, see Fig. 3-centre. Otherwise, we attach to Q a lollipop wrapped in a sphere, as in Fig. 3-right. Of course $c(P, S(Y))$ has the required value and it is easy to show that, with the only exception where Q is a point, so $X = \mathbb{S}^3_{\rm o}$ and $Y = \mathbb{S}^3_{\rm c}(p)$, the spine P constructed is minimal if Q was. A 2-disc is of course a minimal spine of $\mathbb{S}^3_{\rm c}(p)$, and the proof is complete. □

Proposition 2.3. *Let X be a 3-orbifold such that:*

- *X does not contain any bad 2-suborbifold;*
- *X cannot be expressed as an ordinary connected sum between some orbifold and some $\mathbb{S}^3_{\rm c}(p)$.*

Then X admits minimal spines, and any such spine is nuclear.

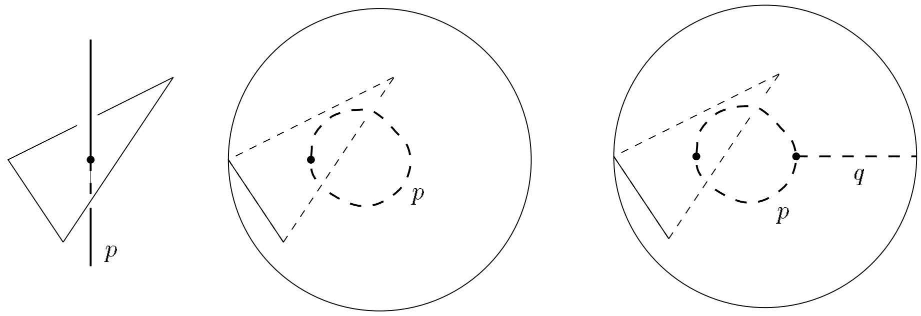

Figure 4: Obstructions to collapsing 2-simplices

Proof. We begin by proving that if we collapse a spine P of X as long as possible we still get a spine of X. To do so, let us triangulate P in such a way that $S(X)$ intersects the 2-simplices only, and each 2-simplex at most once, and let us examine the elementary collapses. Of course we still have a spine of X after collapsing a 1-simplex. The only case where collapsing a 2-simplex σ does not give a spine arises when σ intersects $S(X)$, as shown in Fig. 4-left. Since σ has at least one free edge, the same component B of $X \setminus P$ is incident to both sides of σ. By definition of spine, B is a discal 3-orbifold of ordinary or cyclic type. Depending on which of these cases occurs, we have in X either the situation shown in Fig. 4-centre, which implies that X has an ordinary $\mathbb{S}^3_{\mathrm{c}}(p)$ connected summand, or the situation shown in Fig. 4-right, which implies that X contains the bad 2-suborbifold $(S^2; q)$. A contradiction.

We have shown so far that any spine P of X can be replaced by the *nuclear* spine obtained by collapsing P as long as possible. The conclusion now follows from the obvious remark that X does have spines P with $c(P, S(X)) = c(X)$, and from the following assertions:

1. If P is a nuclear polyhedron and $Q \subset P$ is homeomorphic to P then Q is equal to P;

2. If $\{P_i\}$ is a sequence of nuclear polyhedra with $P_{i+1} \subseteq P_i$ then P_i is eventually constant.

To prove (1), note first that $S(Q) \subset S(P)$. It easily follows that $S(Q) = S(P)$, otherwise Q would be collapsible, but $Q \cong P$. We now claim that the 2-dimensional portion of Q coincides with that of P. By contradiction, let α be a path in the 2-dimensional portion of P which joins a point of

$S(P) = S(Q)$ to a point of $P \setminus Q$. If x is the last point of α which belongs to Q, then x belongs to some face of Q which can be collapsed, so Q is not nuclear, but $Q \cong P$. The proof that Q has the same 1-dimensional portion as P is carried out along the same lines.

Assertion (2) is already implicit in Matveev's statement that minimal spines for manifolds exist, so we will refrain from giving a formal proof. $\square$

Some non-special spines Besides the orbifold $\mathbb{S}^3_*$ already defined above, we will need to consider in the sequel certain orbifolds $(\mathbb{P}^3, F_p)$ and $(L_{3,1}, F_p)$. In both cases F_p is a circle of order p, given by a non-singular fibre of the natural Seifert fibration. To include the case of the manifolds $S^3, \mathbb{P}^3, L_{3,1}$ (those shown by Matveev to be the closed irreducible ones of complexity 0) we stipulate that if K is a knot then K_p denotes K equipped with the cone order p if $p \geqslant 2$, and it denotes the empty set if $p = 1$. Coherently with this choice, we denote $\mathbb{S}^3_{\mathrm{o}}$ also by $\mathbb{S}^3_{\mathrm{c}}(1)$.

Proposition 2.4. $c(\mathbb{S}^3_{\mathrm{c}}(p)) = c(\mathbb{P}^3, F_p) = c(L_{3,1}, F_p) = p - 1$, *and any minimal spine of any of these orbifolds is non-special.*

Proof. The second assertion follows from the first one, because if P is a special spine of (M, K_p), where K is a knot, then P has at least one vertex and it meets K at least once, so $c(P, K_p) \geqslant 1 + (p - 1) > p - 1$. For the same reason it is sufficient to show that for each of the orbifolds in question there is a spine which has no vertices and meets the singular set once. The 2-disc, the projective plane, and the "triple hat"

$$\left\{z \in \mathbb{C} : \ |z| \leqslant 1\right\} \Big/ \left(z \sim w \text{ if } |z| = |w| = 1 \text{ and } z^3 = w^3\right)$$

are such spines for $\mathbb{S}^3_{\mathrm{c}}(p)$, $(\mathbb{P}^3, F_p)$, and $(L_{3,1}, F_p)$, respectively. $\square$

The next result deals with the order-(p, q, r) vertex spherical 3-orbifold $\mathbb{S}^3_{\mathrm{v}}(p, q, r)$ introduced in Section 1.

Proposition 2.5. $c(\mathbb{S}^3_{\mathrm{v}}(p, q, r)) = (p-1)+(q-1)+(r-1)$ *and any minimal spine of this orbifold is not special.*

Proof. Denote the singular set of $\mathbb{S}^3_{\mathrm{v}}(p, q, r)$ by $\theta_{p,q,r}$. A sphere S^2 embedded in S^3 so to separate one vertex of $\theta_{p,q,r}$ from the other one is a spine of the required complexity. The proof is completed along the lines of the previous one. $\square$

Orbifolds with minimal special spines We prove now the main result of the present section:

Theorem 2.6. *If X is an irreducible 3-orbifold and not $\mathbb{S}^3_{\mathrm{c}}(p)$, $(\mathbb{P}^3, F_p)$, $(L_{3,1}, F_p)$, or $\mathbb{S}^3_{\mathrm{v}}(p,q,r)$, then any minimal spine of X is special.*

Proof. Let P be a minimal spine of X. By Proposition 2.3 we can assume that P is nuclear. Then P is the union of an almost-special polyhedron and a graph (see [4]), so it can be non-special only if one of the following occurs:

1. P is a point;
2. P has some purely 1-dimensional portion;
3. P is the triple hat;
4. P is almost special and some 2-dimensional region of P is not a 2-disc.

In case (1) of course X is $\mathbb{S}^3_{\mathrm{o}}$. Suppose we are in case (2), and take a small 2-disc D transversal to an isolated edge of P. The circle ∂D can be seen as a loop on the boundary of some component B of $X \setminus P$. Recall that B is one of the $\mathbb{D}^3_*$'s. The condition that X contains no bad 2-suborbifolds easily implies that ∂D, viewed on ∂B, cannot encircle any of the cone points of ∂B. Therefore ∂D bounds within B a non-singular 2-disc D'. Now $D \cup D'$ is an ordinary sphere, so it bounds an ordinary 3-disc Δ, and by construction $\partial\Delta$ meets P in a single point. It follows that if we dismiss the whole of $P \cap \Delta$ we still have a spine of X. Since P is nuclear and $P \setminus \Delta$ is not, we get a contradiction to the minimality of P.

In case (3) it is easy to check that $|X|$ is necessarily $L_{3,1}$. Moreover every point of intersection between $S(X)$ and P gives two cone points on the boundary of the 3-disc $X \setminus P$, so there is at most one intersection point. It easily follows that X is $(L_{3,1}, F_p)$ for some p.

Let us turn to case (4). One can see that one of the following must occur:

(a) $P = S^2$;

(b) $P = \mathbb{P}^3$;

(c) there exists a 2-dimensional region R of P and a loop γ on R which cuts R into two surfaces neither of which is a 2-disc.

In case (a) we see that X is either $\mathbb{S}^3_{\mathrm{c}}(p)$ or $\mathbb{S}^3_{\mathrm{v}}(p,q,r)$. In case (b) we have $|X| = \mathbb{P}^3$ and, with the same argument used for $L_{3,1}$, we see that

$X = (\mathbb{P}^3, F_p)$. Let us then turn to case (c), and note that γ is orientation-preserving on R, so R is transversely orientable along γ (because $|X|$ is orientable). It follows that there exists an annulus A with core γ which meets P along γ only, and transversely. The boundary components $\gamma_\pm$ of A can now be viewed as loops on components $B_\pm$ of $X \setminus P$, possibly with $B_- = B_+$. We take now within $B_\pm$ a disc $D_\pm$ bounded by $\gamma_\pm$, with minimal possible intersection with $S(X)$, and $D_- \cap D_+$ empty in case $B_- = B_+$. We then get a sphere $\Sigma = D_+ \cup A \cup D_-$. Recall again that on $\partial B_\pm$ there can be zero, two, or three cone points, and note that $\gamma_\pm$ could or not separate one of these cone points from the other one(s). Depending on which case occurs for γ_- and γ_+, the sphere Σ can be either non-singular, or a bad 2-suborbifold, or a spherical 2-suborbifold of type $\mathbb{S}^2_c(p)$. By the irreducibility of X we deduce that the second case is impossible, and in the other cases there is an ordinary or cyclic discal 3-orbifold Δ bounded by Σ. By construction $\partial\Delta \cap P = \gamma$ and $P \cap \Delta$ is not a 2-disc. If Δ is ordinary, we conclude that P is non-minimal replacing $P \cap \Delta$ by a 2-disc bounded by γ, as in [4], because either $P \cap \Delta$ contains vertices or $P \cap \Delta$ properly contains such a 2-disc. Suppose then that Δ is cyclic, with singular set an arc α. By Lemma 2.1, $\alpha \cap P$ consists of at most one point. If $\alpha \cap P$ is empty then we get from P a new spine P' simply by dismissing the whole of $P \cap \Delta$. Similarly, if $\alpha \cap P$ consists of a point, we obtain from P a new spine P' by replacing $P \cap \Delta$ by a 2-disc bounded by γ. Note that in both cases P' has the same intersections as P with $S(X)$. Moreover P' contradicts the minimality of P for precisely the same reasons as in the case where Δ is an ordinary 3-disc. The proof is complete. $\square$

Corollary 2.7. *For every natural n the set*

$$\mathcal{X}_n := \{X \ \textit{irreducible 3-orbifold}: \ c(X) \leqslant n\}$$

is finite.

Duality We define a *triangulation* of a 3-orbifold X to be a triangulation of $|X|$ such that $S(X)$ is a subset of the 1-skeleton. Recall that for manifolds we use a notion of triangulation more flexible than the traditional PL notion (see Section 1). Note also that we do not impose the restriction, used for instance in [2], that each tetrahedron should have at most one singular edge. We now extend Proposition 1.2, showing that for irreducible orbifolds minimal spines typically are dual to triangulations.

Proposition 2.8. *Under the assumptions of Theorem 2.6, dual to any minimal spine of X there is a triangulation.*

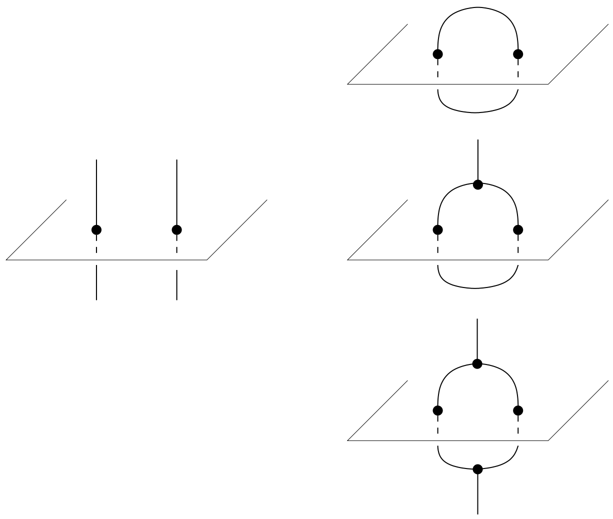

Figure 5: A region of spine meeting the singular locus twice.

Proof. We know that P is special, so dual to it there is a triangulation T of $|X|$, by Proposition 1.2. We now claim that every region R of P meets $S(X)$ at most once. This claim easily yields the conclusion, because we can then isotope the arc of $S(X)$ meeting R, if any, to the edge of T dual to R.

To prove the claim, suppose by contradiction that some region R has $k \geqslant 2$ intersections. Of course $k \leqslant 3$. In case $k = 2$ we see that the two components $B_\pm$ of $X \setminus P$ incident to R must be distinct, because the total number of cone points on their boundaries is at least 4. By definition of spine, $S(B_\pm)$ is either a trivial arc or a trivial Y-graph. This leads to the three possibilities shown in Fig. 5. Correspondingly, one sees either that X has an $\mathbb{S}^3_{\mathrm{c}}(p)$ ordinary connected summand, or that X contains a bad 2-suborbifold, or that X has a cyclic connected summand $\mathbb{S}^3_{\mathrm{v}}(p,q,r)$. All these cases are excluded by the assumptions. Similarly, if $k = 3$ then X has an ordinary connected summand $\mathbb{S}^3_{\mathrm{v}}(p,q,r)$, whence the conclusion. $\square$

3 Spherical splitting of 3-orbifolds

To prove Theorems 0.4 and 0.5 we need to recall more about connected sum of orbifolds than we did in Theorem 0.3. We address the reader to [6] for all proofs and more details.

Splitting systems There is a dual viewpoint to connected sum, which we will mostly employ. To describe it, we first give a definition. If a 3-orbifold Y is bounded by spherical 2-orbifolds, we can canonically associate to Y a closed orbifold $\widehat{Y}$ by attaching the appropriate discal 3-orbifold to each component of ∂Y. We say that $\widehat{Y}$ is obtained by *capping* Y.

Recall now that an operation of connected sum consists of two phases, namely puncturing and gluing. If we consider now a sequence of connected sums, we can always arrange the successive punctures to be disjoint from the spherical 2-orbifolds along which the previous gluings have been performed. If X is the result of the sequence of sums, we then have in X a finite system $\mathcal{S}$ of spherical 2-suborbifolds such that the components of $(X \setminus U(\mathcal{S}))\widehat{\ }$ are the original orbifolds we have summed. (Just as above, U denotes here an open regular neighbourhood, and, for the sake of brevity, we will actually identify $X \setminus U(\mathcal{S})$ to $X \setminus \mathcal{S}$).

Conversely, to any system $\mathcal{S}$ of *separating* spherical 2-orbifolds in a 3-orbifold X there corresponds a realization of X as a connected sum of the components of $(X \setminus \mathcal{S})\widehat{\ }$. We now call *efficient* a finite system of separating spherical 2-suborbifolds which corresponds to an efficient connected sum. To characterize the efficient splitting systems, we call *punctured discal* a 3-orbifold obtained from one of the $\mathbb{D}^3_*$'s by removing a regular neighbourhood of a finite subset.

Proposition 3.1. *A system $\mathcal{S}$ of spherical 2-suborbifolds of a 3-orbifold X is efficient if and only if no component of $X \setminus \mathcal{S}$ is punctured discal, and each component of $(X \setminus \mathcal{S})\widehat{\ }$ is irreducible.*

Uniqueness and stepwise splitting The core of the uniqueness part of Theorem 0.3 is the following result:

Proposition 3.2. *Let $\mathcal{S}$ be an efficient splitting system in a closed 3-orbifold X. Then $\mathcal{S}$ and $(X \setminus \mathcal{S})\widehat{\ }$, as abstract collections of 2- and 3-orbifolds respectively, depend on X only.*

We now state a result which gives a practical recipe to construct an efficient splitting system. Recall that a spherical 2-orbifold is inessential if it bounds a discal 3-orbifold.

Theorem 3.3. *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 $\mathcal{S} = \emptyset$, consider the (non-deterministic) process described by the following steps:*

1. *If all the ordinary spherical* 2*-suborbifolds of* $(X \setminus \mathcal{S})^\frown$ *are inessential, turn to Step 2. Otherwise choose* Σ *as one such* 2*-suborbifold disjoint from* $\mathcal{S}$*, redefine* $\mathcal{S}$ *as the given* $\mathcal{S}$ *union* Σ*, and repeat Step 1;*

2. *If all the cyclic spherical* 2*-suborbifolds of* $(X \setminus \mathcal{S})^\frown$ *are inessential, turn to Step 3. Otherwise choose* Σ *as one such* 2*-suborbifold disjoint from* $\mathcal{S}$*, redefine* $\mathcal{S}$ *as the given* $\mathcal{S}$ *union* Σ*, and repeat Step 2;*

3. *If all the vertex spherical* 2*-suborbifolds of* $(X \setminus \mathcal{S})^\frown$ *are inessential, output* $\mathcal{S}$*. Otherwise choose* Σ *as one such* 2*-suborbifold disjoint from* $\mathcal{S}$*, redefine* $\mathcal{S}$ *as the given* $\mathcal{S}$ *union* Σ*, and repeat Step 3.*

Then the process is finite and the final $\mathcal{S}$ *is an efficient splitting system.*

4 Handles and normal 2-orbifolds

In this section we introduce the notion of handle decomposition for a 3-orbifold, we define normal 2-orbifolds with respect to handle decompositions, and we prove (under suitable assumptions) the fundamental fact that essential spherical 2-orbifolds can be normalized.

Handle decompositions of 3-orbifolds We will use Matveev's terminology of [5], calling *balls*, *beams*, and *plates* respectively the 0-, 1-, and 2-handles of a handle decomposition of a 3-manifold M. In addition, we will call *caps* the 3-handles. Again following Matveev, we will also call *islands* (respectively, *bridges*) the connected components of the attaching loci between the balls and the beams (respectively, the balls and the plates).

If X is a 3-orbifold, we call *orbifold-handle decomposition* of X a handle decomposition of $|X|$ such that the following holds:

- The balls and the beams are disjoint from $S(X)$;
- Each plate $D^2 \times I$ meets $S(X)$ in an arc $\{*\} \times I$;
- Each cap is isomorphic to some $\mathbb{D}^3_*$.

Existence of orbifold-handle decompositions is very easily established.

Normal 2-orbifolds Fix an orbifold-handle decomposition of a 3-orbifold X. We will say that a 2-suborbifold Σ of X is in *normal position* with respect to the decomposition if:

- Σ is disjoint from the caps;
- Σ meets each plate $D^2\times I$ in a family of parallel 2-discs $D^2\times\{t_1,\ldots,t_\nu\}$;
- Σ meets each beam $I \times D^2$ in a set of the form $I \times (\alpha_1 \cup \ldots \cup \alpha_\mu)$, where $\{\alpha_j\}$ is a family of disjoint properly embedded arcs in D^2;
- Σ meets each ball in a union of 2-discs, and the boundary loop of each of these 2-discs passes through each bridge at most once.

We prove now the key result of the section. In the proof we will refer to the "normalization" moves for surfaces described in [5].

Proposition 4.1. *Let X be a closed 3-orbifold which contains no bad 2-suborbifolds and no non-separating spherical suborbifold. Fix an orbifold-handle decomposition of X. Then:*

1. *If X contains some essential ordinary sphere then it contains a normal one;*
2. *Suppose that in X every ordinary sphere is inessential. If X contains an essential cyclic or vertex spherical 2-suborbifold then it contains a normal one of the same type.*

Proof. We prove both points at the same time. Let Σ be the given essential spherical 2-suborbifold. We note that $|\Sigma|$ is a sphere and we apply to it the normalization moves for surfaces, noting that all of them except two are isotopies of Σ (as a 2-suborbifold of X, not just as a surface), so they preserve essentiality. The two exceptions are as follows:

- The compression of $|\Sigma|$ along a 2-disc, followed by the choice of one of the two resulting spheres;
- The move which allows to eliminate double passages along bridges, as shown in Fig. 6 (and explained in its caption).

Since the first move takes place within the union of balls and beams, the compression 2-disc involved does not intersect the singular set. An equivalent way of describing the second move is provided in Fig. 7: we first compress Σ along the 2-disc D and then we dismiss the resulting 2-orbifold Σ', which bounds an ordinary or cyclic discal 3-orbifold, keeping the 2-orbifold Σ''.

According to this discussion, the following claim is sufficient to conclude the proof: *If Σ' and Σ'' are the 2-suborbifolds obtained by compressing Σ*

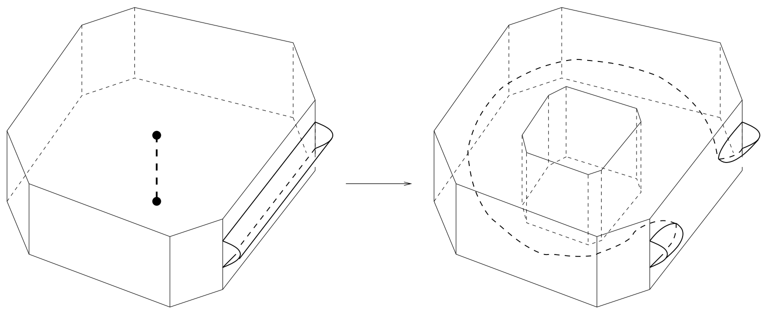

Figure 6: To eliminate a double passage of a surface along a bridge one should first isotope the surface so that it avoids the core of the corresponding plate, and then resize the plate to a neighbourhood of its core.

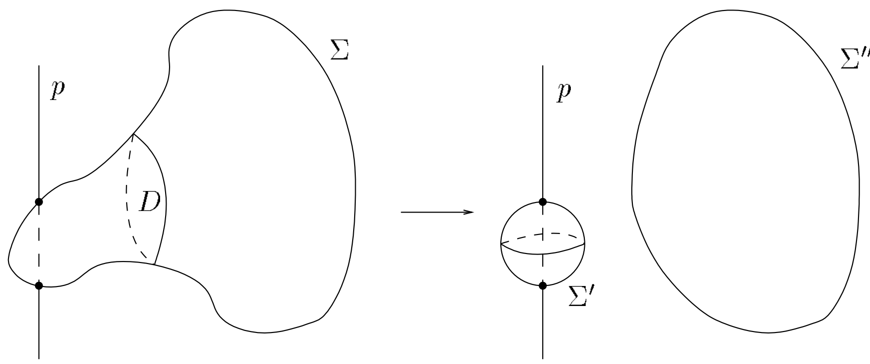

Figure 7: Assuming the core of the plate has order $p \geqslant 1$, the move of Fig. 6 can be seen as the compression shown here followed by the elimination of the inessential ordinary or cyclic spherical 2-orbifold Σ'.

along a 2-disc D which does not meet $S(X)$, then either Σ' or Σ'' is essential and has the same nature as Σ.

Let D' and D'' be the closures of the components of $\Sigma \setminus \partial D$. Since D is non-singular, the assumption that X contains no bad 2-suborbifold easily implies that (up to changing notation) also D'' is non-singular. Therefore $\Sigma' = D \cup D'$ has the same nature as Σ and Σ'' is ordinary, and the conclusion readily follows if we show that either Σ' or Σ'' is essential.

Suppose by contradiction that both Σ' and Σ'' are inessential, *i.e.* that they bound discal 3-orbifolds. Then Σ bounds in X an orbifold of one of the following two types:

- The gluing along a non-singular 2-disc of two discal 3-orbifolds, one of which is ordinary;
- One of the components of a 3-orbifold obtained by cutting a discal 3-orbifold along a non-singular 2-disc.

Any orbifold of such a type is discal, and the proof is complete. □

Normal 2-orbifolds with respect to spines Suppose P is a simple spine of a 3-orbifold X, and consider a triangulation of P, in the traditional PL sense, with the property that 1-simplices are disjoint from $S(X)$ and each 2-simplex meets $S(X)$ in at most one point. Note that the last condition may force subdividing some 2-dimensional portion of P more than a triangulation of P itself would require, because the 2-dimensional components of P can a priori intersect $S(X)$ in as many as three points. Consider the handle decomposition of $|X|$ whose balls, beams, and plates are obtained by thickening the vertices, edges, and triangles of this triangulation of P, and taking as caps the components of what is left. One easily sees that this is also an orbifold-handle decomposition of X. Now let Σ be a normal 2-orbifold with respect to this orbifold-handle decomposition. Then we can attach to each 2-dimensional component of P an integer, corresponding to the number of times Σ runs parallel to the component, and Σ is uniquely determined by the assignment of non-negative integers to the 2-components of P.

5 Additivity under ordinary and cyclic connected sum

In this section we deal with two of the three types of connected sum, leaving the vertex type to the next section. We show that (under some restrictions)

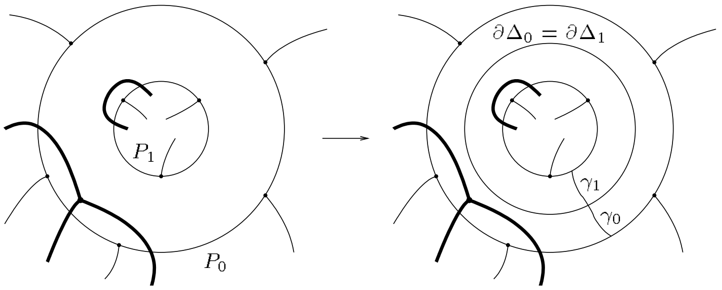

Figure 8: Spine of an ordinary connected sum.

orbifold complexity is additive (in a suitable sense) under these operations.

Subadditivity We begin by giving upper estimates for the complexity of an orbifold obtained as an ordinary or cyclic connected sum of two given orbifolds. For cyclic connected sums it turns out that the estimate depends also on whether the singular components involved in the sum both contain vertices or not. And we will see below that this distinction is unavoidable.

Proposition 5.1. *If X is the ordinary connected sum of X_0 and X_1 then*

$$c(X) \leqslant c(X_0) + c(X_1).$$

Proof. Let P_j be a minimal spine of X_j. To perform the connected sum, we must remove from X_j an ordinary 3-disc Δ_j disjoint from $S(X_j)$, choose a homeomorphism h between $\partial\Delta_0$ and $\partial\Delta_1$, and glue $X_0 \setminus \Delta_0$ to $X_1 \setminus \Delta_1$ along h. Since the position of Δ_j is immaterial, we choose Δ_j to be disjoint from P_j. Therefore $P_0 \sqcup P_1$ can be viewed as a subset of X, and its complement consists of some admissible 3-discs together with one component $S^2 \times (-1,1)$ which contains at most two portions of singular set, separated by $S^2 \times \{0\}$ if there are two of them. The idea is now to construct a spine of X by adding $\big(\{*\} \times (-1,1)\big) \cup \big(S^2 \times \{0\}\big)$ to $P_0 \sqcup P_1$, see Fig. 8

To formalize this idea, we observe that Δ_j is contained in a component of $X_j \setminus P_j$, so we can find within this component an arc γ_j which joins P_j to $\partial\Delta_j$ without meeting $S(X_j)$. We define now $P'_j = P_j \cup \gamma_j \cup \partial\Delta_j$ and we construct P as the gluing of P'_0 and P'_1 under the homeomorphism h between $\partial\Delta_0$ and $\partial\Delta_1$. It is easy to see that P is a spine of X and $c(P, S(X)) = c(P_0, S(X_0)) + c(P_1, S(X_1))$, whence the conclusion. $\square$

Remark 5.2. As opposed to the case of manifolds treated in [4], in the previous proof it was essential to include in the spine P of X the sphere along

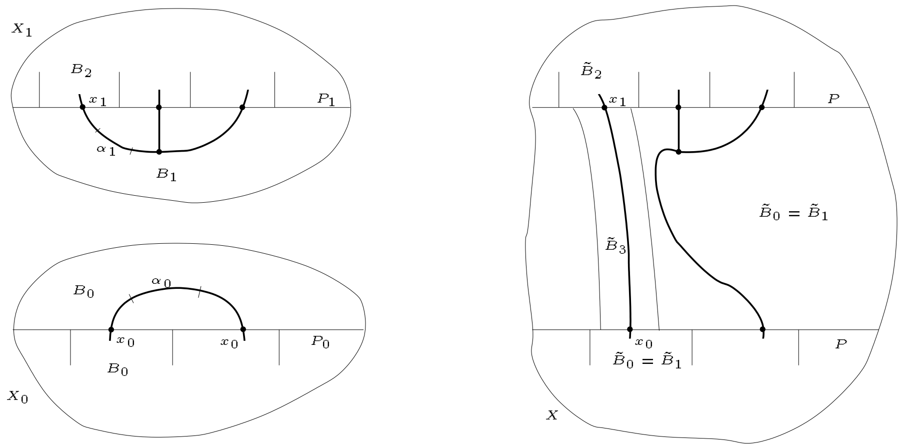

Figure 9: Construction of a spine of a cyclic connected sum.

which the connected sum was performed. This is because the component of $X_j \setminus P_j$ containing Δ_j can be a cyclic or vertex 3-disc. And, if this happens for $j = 0$ and $j = 1$, the gluing of these components along a small boundary 2-disc does not have the required type.

Proposition 5.3. *Let X be a cyclic connected sum of X_0 and X_1 along arcs of order p. Suppose that at least one of the singular components involved in the sum is a knot. Then*

$$c(X) \leqslant c(X_0) + c(X_1) - (p - 1).$$

Proof. Let P_j be a minimal spine of X_j. At the level of the singular set, the connected sum is performed by removing a small arc α_0 from $S(X_0)$ and one α_1 from $S(X_1)$, and by joining together in pairs the ends thus created. Of course we can assume that α_j is disjoint from P_j, so it is contained in a component B_j of $X_j \setminus P_j$. Let $\widetilde{\alpha}_j$ be the component of $S(X_j) \setminus \{\text{vertices}\}$ which contains α_j. By assumption, up to switching indices, $\widetilde{\alpha}_0$ is a knot. Thus, by Lemma 2.1, $\widetilde{\alpha}_0$ is contained in the closure of B_0, and it intersects P_0 in a single point x_0. The same lemma implies that $\widetilde{\alpha}_1$ also meets P_1 in a single point x_1. A schematic representation of the situation is given in Fig. 9-left, assuming that B_1 is a vertex 3-disc. Note that the component B_2 of $X_1 \setminus P_1$ may or not be equal to B_1.

Without loss of generality we can now suppose that the connected sum is performed by removing 3-discs contained in B_0 and B_1. Then $P_0 \sqcup P_1$ can be viewed as a subset of X, and its complement consists of some admissible 3-discs together with a component $S^2\times(-1,1)$. This component contains two portions of singular set, namely an arc which joins $S^2\times\{-1\}$ to $S^2\times\{+1\}$, and another portion which can be an arc or a Y-graph and has precisely one end on $S^2\times\{-1\}$.

We can now create a spine P of X by adding to $P_0\sqcup P_1$ a cylinder in $S^2\times(-1,1)$ which encircles the singular arc, as suggested in Fig. 9-right. Note that $c(P,S(X))=c(P_0,S(X_0))+c(P_1,S(X_1))$. Let us consider now the components of $X\setminus P$, denoted by $\widetilde{B}_j$ in the picture. Then of course $\widetilde{B}_3$ is not equal to $\widetilde{B}_1$ or $\widetilde{B}_2$ (but we could have $\widetilde{B}_1=\widetilde{B}_2$ in case $B_1=B_2$). Therefore we can puncture P near x_0 or near x_1 (but not both, in general), thus reducing by 1 the number of intersection points between P and $S(X)$. The conclusion readily follows. □

Proposition 5.4. *Let X be a cyclic connected sum of X_0 and X_1. Suppose that both the singular components involved in the sum contain vertices. Then*

$$c(X)\leqslant c(X_0)+c(X_1).$$

Proof. In the previous proof, where did we use the assumption that the connected sum should involve at least one singular component without vertices? It was when we supposed that the 3-disc B_0 which contains α_0 intersects $S(X_0)$ in an arc (rather than in a Y-graph). This condition need not be met here, but we can always apply to P_0 a bubble move near a point of intersection between P_0 and $S(X_0)$. This creates a new spine P_0' of X_0 with $c(P_0',S(X_0))=c(P_0,S(X_0))+(p-1)$. We can now apply to P_0' and P_1 precisely the same argument as above, deducing that $c(X)\leqslant c(P_0',S(X_0))+c(P_1,S(X))-(p-1)$, whence the conclusion. □

Superadditivity In this paragraph we prove the estimates opposite to those established in Propositions 5.1, 5.3, and 5.4.

Proposition 5.5. *Let X be a 3-orbifold. Let Σ be a separating ordinary spherical 2-suborbifold of X. Let X_+ and X_- be the components of $(X\setminus\Sigma)\widehat{\ }$. Let P be a spine of X. Suppose that Σ is in normal position with respect to P. Then there exist spines $P_\pm$ of $X_\pm$ such that*

$$c(P_+,S(X_+))+c(P_-,S(X_-))\leqslant c(P,S(X)).$$

Proof. The desired spines are obtained by cutting P along Σ, as in [4]. As we cut we do not create vertices, so $c(P_+)+c(P_-)\leqslant c(P)$. Recall now that a region of P having a positive weight n in Σ gets replaced by n regions, so the intersections between the spine and the singular set could a priori increase after cutting. However, our sphere Σ was supposed to be ordinary, so the regions carrying positive weight do not intersect $S(X)$, and the conclusion easily follows. $\square$

Proposition 5.6. *Let X be a 3-orbifold. Let Σ be an order-p cyclic separating spherical 2-suborbifold of X. Suppose there exists an arc in $S(X)$ joining the two cone points of Σ and not containing vertices of $S(X)$. Let X_+ and X_- be the components of $(X \setminus \Sigma)\hat{}$. Let P be a minimal spine of X. Suppose that Σ is in normal position with respect to P. Then there exist spines $P_\pm$ of $X_\pm$ such that*

$$c(P_+, S(X_+)) + c(P_-, S(X_-)) \leqslant c(P, S(X)) + (p-1).$$

Proof. As in the previous proof, we cut P along $|\Sigma|$, but now we must be careful because the singular set indeed intersects the spine. To be precise, we note that Σ meets $S(X)$ twice, so we have the following possibilities:

(a) There is a region of P which intersects $S(X)$ once and has weight 2 in Σ; all other regions of P having positive weight do not intersect $S(X)$;

(b) There is a region of P which intersects $S(X)$ twice and has weight 1 in Σ; all other regions of P having positive weight do not intersect $S(X)$;

(c) There are two regions regions of P which both intersect $S(X)$ once and have weight 1 in Σ; all other regions of P having positive weight do not intersect $S(X)$.

However we can see that cases (b) and (c) are absurd, because, together with the assumption that the cone points of Σ can be joined in $S(X)$ avoiding vertices, they would imply that there is an arc of $S(X)$ meeting P twice and not containing vertices, which contradicts Lemma 2.1 (note that we have supposed P to be minimal).

The situation is therefore as shown (in a cross-section) in Fig. 10-left, where Σ' and Σ'' denote two of the spherical 2-orbifolds bounding a regular neighbourhood of P (possibly $\Sigma' = \Sigma''$). After cutting we have then the situation of Fig. 10-right, where the new spines are denoted by $\widetilde{P}_\pm$, and $\Sigma_\pm$ are the spherical orbifolds bounding a product neighbourhood of Σ. (To draw the figure we have used the assumption that Σ is separating). We note now that in X_+ we certainly have $\Sigma_+ \neq \Sigma'$, so we can puncture at least one region of $\widetilde{P}_+$, eliminating one intersection point with $S(X_+)$. We define P_+

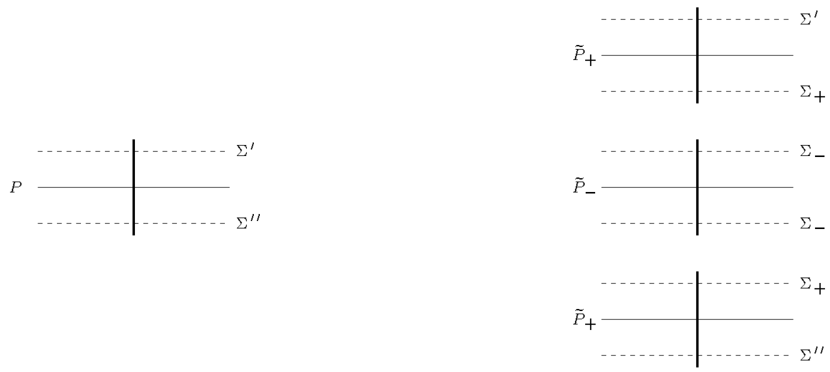

Figure 10: Cutting a spine along a sphere which passes twice through a region.

as the punctured $\widetilde{P}_+$, and P_- as $\widetilde{P}_-$. We note that P_+ is still a spine of X_+ because (by construction) Σ_+ bounds a cyclic discal 3-orbifold in X_+. Since in passing from P to $P_+ \sqcup P_-$ we have not created vertices but we have doubled a point of intersection with the singular locus, the desired estimate readily follows. □

Proposition 5.7. *Let X be a 3-orbifold. Let Σ be an order-p cyclic separating spherical 2-suborbifold of X. Suppose that any arc in $S(X)$ joining the two cone points of Σ contains some vertex of $S(X)$. Let X_+ and X_- be the components of $(X \setminus \Sigma)\widehat{\ }$. Let P be a spine of X. Suppose that Σ is in normal position with respect to P. Then there exist spines $P_\pm$ of $X_\pm$ such that*

$$c(P_+, S(X_+)) + c(P_-, S(X_-)) \leqslant c(P, S(X)).$$

Proof. The proof closely imitates the previous one, except that it is case (a) that is now absurd. We then have the situation of Fig. 11-left, where the four spherical orbifolds $\Sigma_j^{(i)}$ could be anything from equal to each other to distinct from each other. After cutting we then have the situation of Fig. 11-right, to draw which we have again used the fact that Σ separates. Now $\Sigma_1^{(i)}$ is distinct from $\Sigma_\pm$, and $\Sigma_\pm$ bounds a cyclic discal 3-orbifold in $X_\pm$, so we can puncture $\widetilde{P}_\pm$ at least once, getting a new spine $P_\pm$. The desired estimate now follows from the fact that we have created two new intersections with the singular set while cutting P, but we have then eliminated two such intersections by puncturing. □

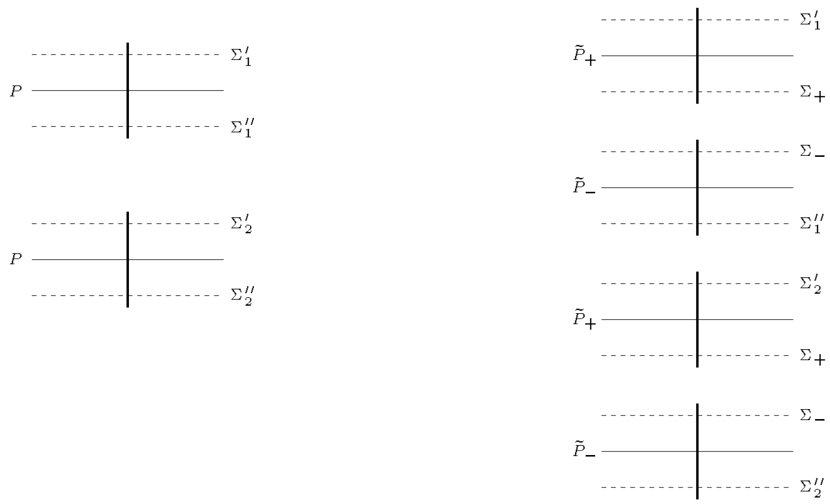

Figure 11: Cutting a spine along a sphere which passes once through two region.

Summarizing statement It is now easy to establish the main result on additivity stated in the Introduction.

Proof of 0.4. By Theorem 3.3, the unique splitting of X given by Theorem 0.3 is obtained by first splitting X along ordinary spheres as long as possible, and then along cyclic spheres. Proposition 4.1 shows that at each step of this successive splitting we can take the sphere to be normal with respect to some handle decomposition. The conclusion then follows from Propositions 5.1 to 5.7. □

6 Estimates on vertex connected sum

When the efficient splitting of an orbifold X into irreducible ones $X_1, \ldots, X_n$ involves cutting along spherical 2-orbifolds of vertex type we cannot prove results as precise as Theorem 0.4. We can however provide upper and lower estimates on the complexity of X in terms of that of $X_1, \ldots, X_n$.

Upper estimates The following result gives an upper bound on the complexity after a vertex connected sum, under assumptions general enough to eventually prove Theorem 0.5.

Proposition 6.1. *Let X_0 and X_1 be orbifolds, and let X be a vertex connected sum of X_0 and X_1. Let P_j be a spine of X_j which is either special or*

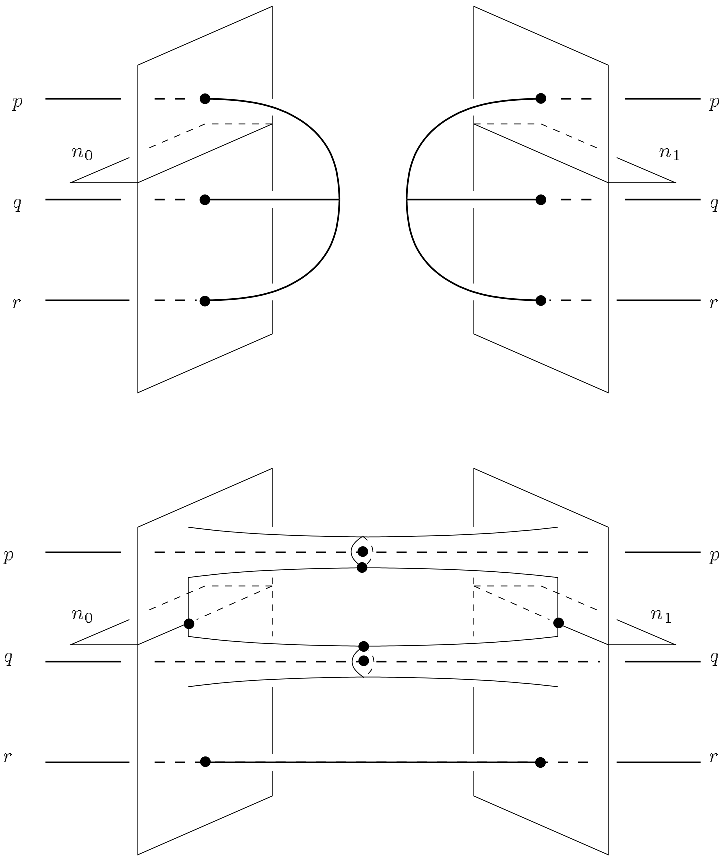

Figure 12: Spine of a vertex connected sum. The region labelled n_j represents n_j parallel regions

homeomorphic to S^2. Then there exists a special spine P of X such that

$$c(P, S(X)) \leqslant 6 \cdot (c(P_0, S(X_0)) + c(P_1, S(X_1))).$$

Proof. Let B_j be the component of $X_j \setminus P_j$ which contains the vertex of $S(X_j)$ along which the connected sum is performed. Let $p_i^{(j)}$, for $i = 1, 2, 3$, be the points of intersection between ∂B_j (which is contained in P_j) and $S(X_j)$. Note that two of these points could coincide, but not all three of them, so we suppose $p_1^{(j)} \neq p_2^{(j)}$. A spine P of X is then constructed as follows (see Fig. 12):

- For $j = 0, 1$, remove from P_j two small 2-discs $D_0^{(j)}$ and $D_1^{(j)}$ around $p_1^{(j)}$ and $p_2^{(j)}$;

- For $i=1,2$, add a tube running from $\partial D_i^{(0)}$ to $\partial D_i^{(1)}$, and a meridinal 2-disc for this tube;
- Add a "rectangular" region with one edge on each tube and, for $j=0,1$, one edge on ∂B_j running from $\partial D_1^{(j)}$ to $\partial D_2^{(j)}$.

The fact that this spine is special is easy to prove and hence left to the reader. Let us discuss how $c(P,S(X))$ relates to $c(P_j,S(X_j))$. Let us first ignore the vertices created on ∂B_0 and ∂B_1 when attaching the rectangle. Then we see that in P there are two more vertices than in $P_0 \sqcup P_1$, and two intersections less with the singular set. So the complexity is not increased at this stage. We are left to give an estimate on how many vertices are created by attaching a region to ∂B_j along a simple path α_j which goes from $p_1^{(j)}$ to $p_2^{(j)}$. Note that vertices arise in two forms:

(A) At the intersections between α_j and $S(P_j)$, as shown in Fig. 12;

(B) At the self-intersections of α_j which can occur within the regions of P_j to which B_j is doubly incident.

To estimate the number of vertices of type (A), we consider a slightly smaller 3-disc B'_j inside B_j, and the trivalent graph $\Gamma_j \subset \partial B'_j$ which projects to $S(X_j)$. Note that this graph has at most $6 \cdot c(P_j)$ edges. Now we can view α_j inside $\partial B'_j$ and arrange it not to touch the edges of Γ_j twice, and not to touch consecutive edges. Therefore the number of vertices of type (A) arising is at most one third the number of edges of Γ_j, hence at most $2 \cdot c(P_j)$.

Turning to type (B), we observe that at most one vertex arises in each region of P_j. Note now that a region touches at least two edges, that at most three regions can be incident to each edge, and that there are $2 \cdot c(P_j)$ edges. Therefore the number of regions is at most $3 \cdot c(P_j)$, and the desired estimate immediately follows. $\square$

Of course in the previous proof the upper estimate could be improved a bit, in particular by separating in $c(P_j,S(X_j))$ the contribution given by the vertices of P_j from that given by the intersections with $S(X_j)$. Since we are only interested in the qualitative fact there exists an estimate which is linear, we refrain from doing this.

Lower estimates We now turn to a lower bound, which as usual employs normal 2-orbifolds.

Proposition 6.2. *Let X be a 3-orbifold as in Theorem 0.3. Let Σ be a separating vertex spherical 2-suborbifold X. Let X_+ and X_- be the components of $(X \setminus \Sigma)\widehat{\ }$. Let P be a spine of X. Suppose that Σ is in normal position with respect to P. Then there exist spines $P_\pm$ of $X_\pm$ such that*

$$c(P_+, S(X_+)) + c(P_-, S(X_-)) \leqslant 4c(P, S(X)).$$

Proof. Let $\widetilde{P}_\pm$ be obtained by cutting P along Σ, as above. Of course $\widetilde{P}_\pm$ is a spine of $X_\pm$, and the operation of cutting along Σ obviously does not increase the total number of vertices. This operation increases by 3 the number of intersections with the singular set, whence the conclusion at once. $\square$

As for the upper estimates, we have not tried to give in the previous proposition an optimal lower one.

Summarizing result We can now prove the estimation result stated in the introduction.

Proof of 0.5. We begin by the upper estimate. Up to reordering we can assume that the connected sums between $X_1, \ldots, X_k$ are of vertex type, and the other ones are not. Then $X_1, \ldots, X_k$ are different from $\mathbb{S}^3_{\rm c}$, $(\mathbb{P}^3, F_p)$, and $(L_{3,1}, F_p)$, so by Theorem 2.6 they have special polyhedra or spheres as minimal spines. Then Proposition 6.1 implies that

$$c(X_1\# \ldots \# X_k) \leqslant 6^{k-1}\big(c(X_1) + \ldots c(X_k)\big)$$

and the conclusion easily follows from Propositions 5.1, 5.3, and 5.4.

The lower estimate immediately follows from Theorems 0.3 and 3.3, and Propositions 4.1, 5.5, 5.6, 5.7, and 6.2. $\square$

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

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

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

[6] C. Petronio, *Spherical splitting of 3-orbifolds*, `math.GT/0409606`.

[7] C. Rourke – B. Sanderson, "Introduction to piecewise-linear topology," Ergebn. der Math. Vol. 69, Springer-Verlag, New York-Heidelberg, 1972.

[8] 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