Journal of Topology 1 (2008) 703–745
doi:10.1112/jtopol/jtn017

3-manifolds efficiently bound 4-manifolds

Francesco Costantino and Dylan Thurston

Abstract

It has been known since 1954 that every 3-manifold bounds a 4-manifold. Thus, for instance, every 3-manifold has a surgery diagram. There are several proofs of this fact, but little attention has been paid to the complexity of the 4-manifold produced. Given a 3-manifold M^3 of complexity n, we construct a 4-manifold bounded by M of complexity $O(n^2)$, where the 'complexity' of a piecewise-linear manifold is the minimum number of n-simplices in a triangulation.

The proof goes through the notion of 'shadow complexity' of a 3-manifold M. A shadow of M is a well-behaved 2-dimensional spine of a 4-manifold bounded by M. We further prove that, for a manifold M satisfying the geometrization conjecture with Gromov norm G and shadow complexity S, we have $c_1 G \leqslant S \leqslant c_2 G^2$, for suitable constants c_1, c_2. In particular, the manifolds with shadow complexity 0 are the graph manifolds.

In addition, we give an $O(n^4)$ bound for the complexity of a spin 4-manifold bounding a given spin 3-manifold. We also show that every stable map from a 3-manifold M with Gromov norm G to $\mathbb{R}^2$ has at least $G/10$ crossing singularities, and if M is hyperbolic there is a map with at most $c_3 G^2$ crossing singularities.

1. *Introduction*

Among the different methods of representing 3-manifolds combinatorially, two of the most popular are triangulations and surgery on a link. A triangulation is a very natural way to represent 3-manifolds, and any other representation of a 3-manifold is easy to turn into a triangulation. On the other hand, although some 3-manifold invariants may be computed directly from a triangulation (for example, the Turaev–Viro invariants), not all can be, and it is difficult to visualize the combinatorial structure of a triangulation.

A more typical way to present a 3-manifold is via Dehn surgery on a link. In practice, there are simple descriptions of small manifolds via surgery, and this is generally the preferred way of representing manifolds. There are many more invariants that may be computed directly from a surgery diagram, like the Witten–Reshetikhin–Turaev invariants. It is easy to turn a surgery diagram into a triangulation of the manifold [40]. However, for the other direction, passing from triangulations to surgery diagrams, little seems to be known. In particular, it is an open question whether a surgery diagram must (asymptotically) be more complicated than a triangulation. For a more general setting of this problem, consider that if all the surgery coefficients are integers, a surgery diagram naturally gives a 4-manifold bounded by the 3-manifold. This leads us to ask the central question of the paper.

Question 1.1. How efficiently do 3-manifolds bound 4-manifolds?

Received 13 November 2007.

2000 *Mathematics Subject Classification* 57N70 (primary), 57M20, 57R45, 57R15 (secondary).

The first author was supported by a Marie Curie Fellowship issued by the European Community and hosted by Institut de Recherche Mathémathique Avancée de Strasbourg. The second author was supported by NSF Grant DMS-0071550, Harvard University, the University of Pisa, and a Sloan Research Fellowship.

To make this question more precise, let us make some definitions.

DEFINITION 1.2. A Δ-*complex* is the quotient of a disjoint union of simplices by identifications of their faces. (See [13, Section 2.1] for a complete definition. These are also semi-simplicial complexes in the sense of Eilenberg and Zilber [6].) A Δ-triangulation is a Δ-complex whose underlying topological space is a manifold.

DEFINITION 1.3. The *complexity* of a piecewise-linear oriented n-manifold M^n is the minimal number of n-simplices in a Δ-triangulation of M,

$$C(M^n) = \min_{\text{Triang. } T \text{ of } M} \text{no. of } n\text{-simplices in } T. \tag{1.1}$$

REMARK 1.4. Since the second barycentric subdivision of a Δ-triangulation is an ordinary simplicial triangulation, $C(M)$ would change only by at most a constant factor if we insisted on simplicial triangulations in Definition 1.3.

DEFINITION 1.5. The 3-*dimensional boundary complexity function* $G_3(k)$ is the minimal complexity such that every 3-manifold of complexity at most k is bounded by a 4-manifold of complexity at most $G_3(k)$.

We can think of $G_3(k)$ as a kind of topological isoperimetric inequality. We can now give a concrete version of our original Question 1.1.

QUESTION 1.6. What is the asymptotic growth rate of G_3?

The first main result of this paper is that $G_3(k) = O(k^2)$. More precisely, we have the following theorem, which appears in Section 5.

THEOREM 5.2. *If a 3-manifold M has a Δ-triangulation with t tetrahedra, then there exists a 4-manifold W such that $\partial W = M$ and W has a Δ-triangulation with $O(t^2)$ simplices. Moreover, W has 'bounded geometry'. That is, there exists an integer c (not depending on M and W) such that each vertex of the triangulation of W is contained in fewer than c simplices.*

The fact that W has bounded geometry makes the resulting representation nicer; in particular, to check whether the topological space resulting from a triangulation is a manifold, you need to decide whether the link of each simplex is a sphere. This is easy for dimension $n \leqslant 3$, in NP for $n = 4$ (see [34]), unknown for $n = 5$, and undecidable for dimension $n > 5$ (see [20, 24, 25]). In all cases, such complexity issues do not arise if the triangulation has bounded geometry. Note that there is an evident linear lower bound for $G_3(k)$, since a triangulation for a 4-manifold also gives a triangulation of its boundary.

We also prove a number of other related bounds that do not directly refer to 4-manifolds. For instance, we have the following bound in terms of surgery, proved in Section 5.

THEOREM 5.6. *A finite-volume hyperbolic 3-manifold with volume V has a rational surgery diagram with $O(V^2)$ crossings.*

Note that there may be an infinite number of 3-manifolds with volume less than the bound V, and likewise an infinite number of surgeries on a given link diagram; but in both cases the manifolds come in families with some structure. In this case as well, there is a linear lower bound: there are at least V/v_{oct} crossings in any surgery diagram for M, where $v_{\text{oct}} \approx 3.66$ is the volume of a regular ideal hyperbolic octahedron. A somewhat weaker lower bound was

proved by Lackenby [19]; the bound using v_{oct} comes from a decomposition into ideal octahedra, one per crossing [29, 35].

For a clean statement about general 3-manifolds, we use the crucial notion of *shadows*, which we recall in Section 3. For now, we just need to know that shadows are certain kinds of decorated 2-complexes that can be used to represent both a 4-manifold and a 3-manifold (on the boundary of the 4-manifold), and that a coarse notion of the complexity of a shadow is the number of its *vertices*. There are an infinite number of 3-manifolds with shadows with a given number of vertices, but as with hyperbolic volume and surgeries on a given link, they come in families that can be understood. The *shadow complexity* $\text{sc}(M)$ of a 3-manifold M is the minimal number of vertices in any shadow for M.

The following theorems combine to say that the shadow complexity gives a coarse estimate of the hyperbolic volume.

THEOREMS 3.37 and 5.5. *There is a universal constant $C > 0$ such that every geometric 3-manifold M, with boundary empty or a union of tori, satisfies*

$$\frac{v_{\text{tet}}}{2v_{\text{oct}}}\|M\| \leqslant \text{sc}(M) \leqslant C\|M\|^2.$$

The lower bound on $\text{sc}(M)$ holds for all 3-manifolds.

Here $v_{\text{tet}} \approx 1.01$ is the volume of a regular ideal hyperbolic tetrahedron and v_{oct} is as above. A *geometric manifold* is one that satisfies the geometrization conjecture [36]: it can be cut along spheres and tori into pieces admitting a geometric structure. $\|M\|$ is the Gromov norm of M, which is defined for any 3-manifold, and for a geometric 3-manifold is $1/v_{\text{tet}}$ times the sum of the volumes of the hyperbolic pieces.

Note that there is no constant term in these theorems. The manifolds with shadows with no vertices are the *graph manifolds*, the geometric manifolds with no hyperbolic pieces (see Proposition 3.31).

Our techniques are based on maps from 3-manifolds to surfaces, so we can also phrase the bounds in terms of the singularities of such maps. A *crossing singularity* is a singularity of the type that we consider in Section 4.4: a point in $\mathbb{R}^2$ with two indefinite fold points in its inverse image. For more background on the classification of the stable singularities of a map from a 3-manifold to a 2-manifold, see Levine [22, 23].

THEOREMS 3.38 and 5.7. *A 3-manifold M has at least $\|M\|/10$ crossing singularities in any smooth, stable map $\pi : M \to \mathbb{R}^2$. There is a universal constant C such that if M is hyperbolic, then M has a map to $\mathbb{R}^2$ with $C\|M\|^2$ crossing singularities.*

One related theorem was previously known: Saeki [33] showed that the manifolds with maps to a surface with no crossing singularities are the graph manifolds. Gromov independently proved [8] both these upper and lower bounds after this paper was written, as explained in Section 1.3.

We also prove bounds for the complexity for 3-manifolds to bound special types of 4-manifolds in the following theorems.

THEOREM 5.3. *A 3-manifold with a triangulation with k tetrahedra is the boundary of a simply connected 4-manifold with $O(k^2)$ 4-simplices.*

THEOREM 6.1. *A 3-manifold with a triangulation with k tetrahedra is the boundary of a spin 4-manifold with $O(k^4)$ 4-simplices.*

All the constants in these theorems can be made explicit, but since in general they are quite bad, we have not given them explicitly. The exceptions are those in Theorems 3.37 and 3.38, which are the best possible.

As one application of the results above, let us mention computing invariants of 3-manifolds. There are a number of 3-manifold invariants that are most easily computed from a 4-manifold with boundary. (Often this is done via surgery diagrams, so the 4-manifold is simply connected, but there are usually more general constructions as well.) For instance, the Witten–Reshitikhin–Turaev (WRT) quantum invariants are of this form [30, 31, 39], as is the Casson invariant [21]. (The original definition of the Casson invariant is 3-dimensional, but to compute it in practice, the surgery formula is much easier.)

As one concrete example, Kirby and Melvin explained [17, 18] how to compute the WRT invariant at a fourth root of unity as a sum over spin structures, which can be done concretely given a surgery diagram. Although they show that the exact evaluation is NP-hard, our results imply that the sum can be approximated (up to some error) in polynomial time using random sampling over spin structures: for any given spin structure, we can, in polynomial time, find a 4-manifold which spin bounds the given 3-manifold and we can therefore compute the summand at this spin structure. This contrasts with an unpublished result of Kitaev and Bravyi, who showed that computing (up to the same error) the partition function of the corresponding 2-dimensional TQFT is BQP-complete, as soon as we allow evaluation of observables on closed curves.

1.1. *Plan of the paper*

In the remainder of the introductory section, we survey some related work, first on different kinds of topological isoperimetric functions, and second on other work considering our main tool, namely stable maps from a 3-manifold to $\mathbb{R}^2$. In Section 2, we sketch our construction in the smooth setting and introduce the crucial tool of the Stein factorization, which shows how 2-complexes naturally arise. This section is not logically necessary for the rest of the paper, although it does provide helpful motivation and a guide to the proof. These 2-complexes that arise are studied more abstractly in Section 3, where we review the definition of shadow surfaces and prove a number of their properties; here we also use the Gromov norm to prove all the lower bounds in the theorems above. In Section 4 we give our main tool, a construction of a shadow from a triangulated 3-manifold with a map to the plane, together with a bound on the complexity of the resulting shadow. Section 4 is independent of Section 3 except for the definition of shadows, and only uses Section 2 as motivation, so the impatient reader can skip to there. In Section 5, we use this construction to prove the upper bounds of our main theorems (except for the spin-bound case, Theorem 6.1) and see precisely how shadow complexity relates to geometric notions on the complexity of the manifold. Finally, in Section 6 we show how to modify an arbitrary shadow to get a 4-manifold that spin bounds a specified spin structure on a 3-manifold, while controlling the complexity.

1.2. *Related questions*

Although the question we consider does not seem to have been previously addressed, there has been related work. Perhaps the closest is the work on distance in the pants complex and hyperbolic volumes. The pants complex is closely related to shadows; in particular, a sequence of moves of length n in the pants complex can be turned into a shadow with n vertices for a 3-manifold which has two boundary components, so that the natural pants decomposition of the boundary components corresponds to the start and end of the sequence of moves.

THEOREM 1.7 (Brock [1, 2]). *Given a surface S of genus $g \geqslant 2$, there are constants C_1, C_2 so that for every pseudo-Anosov map $\psi : S \to S$, we have*

$$C_1 \|\psi\|_{\mathrm{Pants}} \leqslant \mathrm{vol}(T_\psi) \leqslant C_2 \|\psi\|_{\mathrm{Pants}},$$

where T_ψ is the mapping torus of ψ and $\|\psi\|_{\mathrm{Pants}}$ is the translation distance in the pants complex.

By the relation between moves in the pants complex and shadows mentioned above, this shows that for 3-manifolds that fiber over the circle with fiber a surface of fixed genus, shadow complexity is bounded above and below by a linear function of the hyperbolic volume. However, the constant depends on the genus in an uncontrolled way. Our result gives a quadratic bound, but with an explicit constant not depending on the genus. Brock's construction also produces shadows (and 4-manifolds) of a particular type.

More recently, Brock and Souto have announced that there is a similar bound for manifolds with a Heegaard splitting with a fixed genus. In our language, their result says that a hyperbolic manifold with a strongly irreducible Heegaard splitting of genus g has a shadow diagram where the number of vertices is bounded by a linear function of the volume, with a constant of proportionality depending only on the genus. (The result is probably true without the assumption that the Heegaard splitting is strongly irreducible, but the statement becomes more delicate in the language of the pants complex and we have not checked the details.) Their method of proof does not produce any explicit constants.

There has also been work on the question of polygonal curves in $\mathbb{R}^3$ bounding surfaces.

DEFINITION 1.8. The *surface isoperimetric function* $G_{\text{surf}}(k)$ is the minimal number, such that every closed polygonal curve γ in $\mathbb{R}^3$ with at most k segments bounds an oriented polygonal surface Σ with at most $G_{\text{surf}}(k)$ triangles.

THEOREM 1.9 (Hass–Lagarias [9]). *We have* $\frac{1}{2}k^2 \leqslant G_{\text{surf}}(k) \leqslant 7k^2$.

This result contrasts sharply with the situation when we ask for the spanning surface Σ to be a disk.

DEFINITION 1.10. The *disk isoperimetric function* $G_{\text{disk}}(k)$ is the minimal number such that every closed polygonal curve γ in $\mathbb{R}^3$ with at most k segments bounds an oriented polygonal disk D with at most $G_{\text{surf}}(k)$ triangles.

THEOREM 1.11 (Hass–Snoeyink–Thurston [11]). *We have* $G_{\text{disk}}(k) = e^{\Omega(k)}$. *That is, there is a constant* C *such that, for sufficiently large* k, $G_{\text{disk}}(k) \geqslant e^{Ck}$.

THEOREM 1.12 (Hass–Lagarias–Thurston [10]). *We have* $G_{disk}(k) = e^{O(k^2)}$. *That is, there is a constant* C *such that, for sufficiently large* k, $G_{\text{disk}}(k) \leqslant e^{Ck^2}$.

Although there is a large gap between these upper and lower bounds, both bounds are substantially larger than the bounds in Theorem 1.9, which was about arbitrary oriented surfaces.

There is an analogous question on the growth of G_{disk} for 3-manifolds rather than curves: namely, asking for 4-balls bounding a 3-sphere with a given triangulation on the boundary. As stated, this is not an interesting question, since we can construct such a triangulation by taking the triangulated 3-ball and coning it to a point. This is related to the somewhat unsatisfactory nature of the 4-manifold complexity (mentioned earlier). A more interesting question might involve 4-manifold triangulations where the vertices have bounded geometry. For a somewhat different question, there are known upper bounds.

DEFINITION 1.13. The *Pachner isoperimetric function* $G_{\text{Pachner}}(k)$ is the maximum over all triangulations T of the 3-sphere with at most k simplices of the minimum number of Pachner moves required to relate T to the standard triangulation, the boundary of a 4-simplex.

THEOREM 1.14 (King [16], Mijatović [28]). *We have* $G_{\text{Pachner}}(k) = e^{O(k^2)}$.

Note that a sequence of Pachner moves as in the definition gives you, in particular, a triangulation of the 4-ball, although you only get very special triangulations of the 4-ball in this way.

As in the case of polygonal surfaces and disks, this upper bound is much larger than the polynomial bound that we obtain. King [16] also constructs triangulations of S^3 which seem likely to require a large number of Pachner moves to simplify.

1.3. *Other work*

The central construction in our proof, a generic smooth map from a 3-manifold to $\mathbb{R}^2$, has been considered by several previous authors, sometimes with little contact with each other. These maps were probably first considered by Burlet and de Rham [3], who showed that the 3-manifolds admitting a map with only definite fold singularities are connected sums of $S^1 \times S^2$ (including S^3). They also introduced the Stein factorization. Levine [23] clarified the structure of the singularities and studied, for instance, related immersions of the 3-manifold into $\mathbb{R}^4$. Burlet and de Rham's result was extended by Saeki [33], who showed that the 3-manifolds admitting a map without codimension-2 singularities (that is, only definite or indefinite folds) are the graph manifolds.

Rubinstein and Scharlemann [32] constructed a map from the complement of a graph in a 3-manifold to $\mathbb{R}^2$ from a pair of Heegaard splittings and used this to bound the number of stabilizations required to turn one splitting into the other. Much of the analysis is similar to ours.

Also independently, Hatcher and Thurston [14] considered Morse functions on a orientable surface to show that its mapping class group is finitely presented. To get a set of generators, they considered one-parameter deformations of the Morse function. Note that a one-parameter family of maps from Σ to $\mathbb{R}$ is a map from the 3-manifold $\Sigma \times [0,1]$ to $\mathbb{R}^2$.

In a slightly different context, Hatcher's proof of the Smale conjecture [12], that the space of smooth 2-spheres in $\mathbb{R}^3$ is contractible, uses the Stein factorization of a map from S^2 to $\mathbb{R}^2$. Hong, McCullough, and Rubinstein recently combined this approach with the Rubinstein–Scharlemann techniques in their proof of the Smale conjecture for lens spaces [15].

On the other side of the story, Turaev [37, 38, 39] introduced shadow surfaces as the most natural objects on which the Reshetikhin–Turaev quantum invariants are defined. He observed that you could construct both a 3-manifold and a 4-manifold with boundary the 3-manifold from a shadow surface.

Most recently, after this paper was first written, we became aware of the independent work of Gromov [8], who studies questions on the required singularities of smooth maps in much greater generality. In particular, he proves the lower bound on crossing number (Theorem 3.38) without the precise constant [8, Section 3], sketches a proof of the upper bound on crossing number (Theorem 5.7) in the case that the injectivity radius is bounded, and proves an optimality result for the quadratic upper bound on maps.

THEOREM 1.15 (Gromov [8, Section 6.3]). *For every closed hyperbolic 3-manifold M, there are a constant $C > 0$ and an infinite sequence $\{M_i\}$ of finite coverings of degree d_i of M, such that for every surface S of genus g and maps $\pi_i : M_i \to S$, the number of crossing singularities of π_i is at least $Cd_i^2/(g+1)$.*

Thus, several authors have been considering nearly the same objects (shadow surfaces on the one hand and the Stein factorization of a map from M^3 to $\mathbb{R}^2$ on the other hand) for several years. The gleams are key topological data from the shadow surface point of view, since

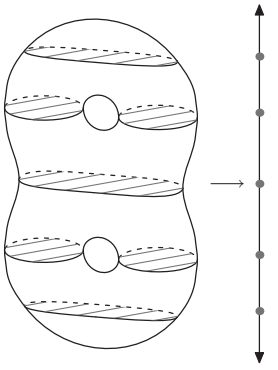

FIGURE 1. *Proving that every surface bounds a 3-manifold: a generic map of a surface to* $\mathbb{R}$*, with regular values marked and disks glued into the inverse image of the regular values.*

they let you reconstruct the 3-manifold, but they were not explicitly described by the authors writing on Stein factorizations, although it is implicitly present.

Acknowledgements. We would like to thank Riccardo Benedetti, Mikhail Gromov, Simon King, Robion Kirby, William Thurston, Vladimir Turaev, and an anonymous referee most warmly for their encouraging comments and suggestions.

2. 4-*manifolds from stable maps*

In order to prove that 3-manifolds efficiently bound 4-manifolds, we start by sketching a proof that 3-manifolds do bound 4-manifolds. In Section 4, we will analyze a version of the proof in the PL setting and give a bound on the complexity of the resulting 4-manifold.

Consider an oriented, smooth, closed 3-manifold M^3 and a generic smooth map f from M to $\mathbb{R}^2$. At a regular value $x \in \mathbb{R}^2$, the inverse image $f^{-1}(x)$ consists of an oriented union of circles. To construct a 4-manifold, we glue a disk to each of these circles away from critical values and then extend across the singularities in codimension 1 and codimension 2.

2.1. *Pants decompositions from Morse functions*

To get some idea of what the singularities look like, we first do the analysis of extension across singularities one dimension down: let us prove that every oriented 2-manifold Σ^2 bounds a 3-manifold. Consider a generic smooth map f from Σ to $\mathbb{R}$, that is, a Morse function. The inverse image of a regular value is again a union of circles. Glue in disks to each of these circles as in Figure 1. More properly, take $\Sigma \times [0,1]$, pick a regular value in each component of $\mathbb{R}$ minus the singular set, and attach 2-handles along the circles appearing in the inverse image of the chosen regular values. The result is a 3-manifold with one boundary component which is Σ and other boundary components corresponding to the singular values of f.

The singular values of a Morse function, locally in the domain Σ, are well known: they are critical points with a quadratic form which is definite (index 0 or 2, minima or maxima) or indefinite (index 1, saddle points). Since our construction works with the entire inverse image of a regular value, we need to understand the singularities locally in the range $\mathbb{R}$; that is, we need to know the connected components of inverse images of a critical value. This is easy for the definite singularities.

Let $p_0 \in \Sigma$ be a saddle point, and let $x_0 \in \mathbb{R}$ be its image. Near p_0, $f^{-1}(x_0)$ is a cross. For x above and below x_0, $f^{-1}(x)$ is locally the cross which is smoothed out in the two possible ways, as on the top of Figure 2. Note that the orientations of Σ and $\mathbb{R}$ induce an orientation of $f^{-1}(x)$

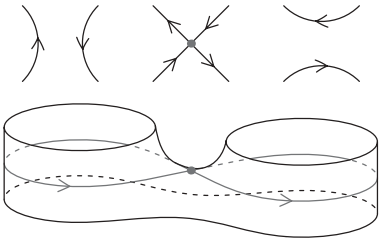

FIGURE 2. *Analyzing the saddle singularity.*

for all $x \in \mathbb{R}$ except at critical points in Σ, so both of these smoothings must be oriented, so the arms of the cross must be oriented alternately in and out. The connected component of $f^{-1}(x_0)$ containing p_0 must join the arms of the cross in an orientation-preserving way and is therefore a figure-8 graph ⧓. This implies that for a small interval I containing x_0, $f^{-1}(I)$ is a pair of pants, as illustrated at the bottom of Figure 2.

To finish constructing the 3-manifold, recall that in the previous step we glued in disks at all the regular values. Near this pair of pants, this means that we have closed off each hole in the pair of pants, and the boundary component that we are trying to fill in is just a sphere, which we can fill in with a ball.

An easier analysis shows that the surface that we need to fill in the other cases of a maximum or a minimum is again a sphere.

Notice where the proof breaks down if we do not assume that Σ is oriented: there is then another possibility for the inverse image of the critical value, with opposite arms of the cross attached to each other: ⟲. In this case, the surface that we are left to fill in turns out to be $\mathbb{RP}^2$, which does not bound a 3-manifold.

In a similar way, we can analyze the possible stable singularities of a smooth map from a 3-manifold M to $\mathbb{R}^2$. We glue in a disk (a 2-handle) to each circle in the inverse image of a regular point, extend across codimension-1 singularities by attaching 3-handles (the singularities look just like the singularities we analyzed for the case of a surface, crossed with $\mathbb{R}$), and then consider the codimension-2 singularities. In Section 4.4, we will analyze the codimension-2 singularities (in the PL category) and see that the remaining boundary from each codimension-2 singularity is S^3, which can be filled in by attaching a 4-handle.

2.2. *Stein factorization and shadow surfaces*

For a more global view, we can consider the *Stein factorization* $f = g \circ h$ of these maps. The Stein factorization of a map f with compact fibers decomposes it as the composition of a map h with connected fibers and a map g which is finite-to-one. That is, h is the quotient onto the space of connected components of the fibers of f. See Figure 3 for an example.

For a stable map from an oriented surface Σ to $\mathbb{R}$, the Stein factorization is generically a 1-manifold, with singularities from the critical points. Concretely, it is a graph with vertices which have valence 1 (at definite singularities) or valence 3 (at indefinite singularities). The surface is a circle bundle over this Stein graph Γ at generic points. Likewise, the 3-manifold that we constructed (with boundary Σ) is a disk bundle over Γ at generic points. In fact, the 3-manifold collapses onto Γ. If we collapse all the valence 1 ends, we may think of Γ as representing a pair-of-pants decomposition of Σ; each circle in the pair-of-pants decomposition bounds a disk in the 3-manifold.

For a stable map from a 3-manifold to $\mathbb{R}^2$, on the other hand, the Stein factorization is generically a surface. The codimension-1 singularities of the Stein surface are products of the lower-dimensional singularities with an interval, and have one or three sheets meeting at an

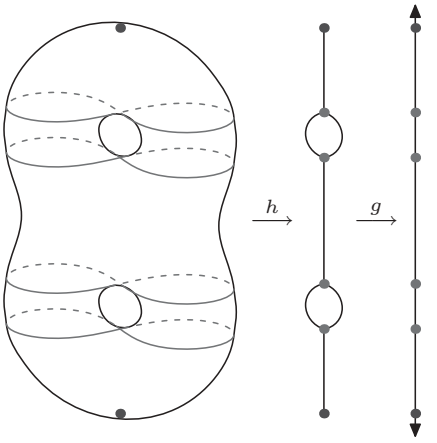

FIGURE 3. *The Stein factorization $f = g \circ h$ of the map in Figure 1.*

FIGURE 4. *Some local models for the Stein factorization of a map from a 3-manifold to $\mathbb{R}^2$ in codimension 0, 1, and 2. In each picture, the map to the plane is the vertical projection.*

edge at what we will call definite or indefinite folds, respectively. In codimension 2, there are a few different configurations of how the surface can meet, the most interesting of which are shown in Figure 4.

The 3-manifold is a circle bundle over the Stein surface at generic points and the 4-manifold is generically a disk bundle. As in the previous case, it turns out that the 4-manifold collapses onto the Stein surface. The resulting surface is very close to a shadow surface.

Unlike in the lower-dimensional case, the surface does not determine the 4-manifold (or the 3-manifold), even after you fix a standard local model of how the surface sits inside the 4-manifold. The additional data that you need are the *gleams*, numbers associated to the 2-dimensional regions of the surface; see Definition 3.7.

3. *Shadow surfaces*

We will now define shadow surfaces and shadows of 3-manifolds and give a few examples. In Section 3.2, we will extend the definitions to 3-manifolds with boundary and an embedded, framed graph. For a more detailed though introductory account of shadows of 3- and 4-manifolds, see [**5**]. Note that these are slightly different from the Stein surfaces mentioned in Section 2.2. We prefer shadow surfaces as the fundamental object, since they are a little more symmetric and regular than Stein surfaces. From now on every manifold will be PL compact and oriented unless explicitly stated otherwise, and every polyhedron will be finite; we also recall that, in dimensions 3 and 4, each PL manifold has a unique smooth structure and vice versa.

3.1. *Shadows of 3-manifolds*

For simplicity, we will first define shadows in the case when there is no boundary, appropriate for 3-manifolds without boundary or other decorations; in the next section we will extend this.

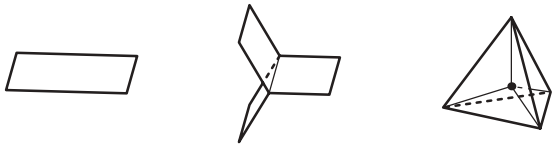

FIGURE 5. *For each point of a simple polyhedron embedded in a 4-manifold, there is a local chart (U, ϕ) in which the polyhedron is flat in the sense that is embedded in a 3-dimensional plane, and in this plane it appears as in one of the three models shown above.*

DEFINITION 3.1. A simple polyhedron P is a compact topological space where every point has a neighborhood homeomorphic to a neighborhood of the central point in one of the local models depicted in Figure 5, namely $\mathbb{R}^2$, $\mathbb{R} \times \text{Cone(three points)}$, and Cone(1-skeleton of tetrahedron). The set of points without a local model of the leftmost type forms a 4-valent graph, called the *singular set* of the polyhedron and denoted $\text{Sing}(P)$. The vertices of $\text{Sing}(P)$ are called *vertices* of P. The connected components of $P \setminus \text{Sing}(P)$ are called the *regions* of P. The set of points of P whose local models correspond to the boundaries of the blocks shown in the figure is called the *boundary* of P and is denoted ∂P; P is said to be *closed* if it has an empty boundary. A region is *internal* if its closure does not touch ∂P. A simple polyhedron is *standard* if every region of P is a disk, and hence $\text{Sing}(P)$ has no circle components.

DEFINITION 3.2. Let W be a PL, compact, and oriented 4-manifold. $P \subset W$ is a *shadow* for W if P is a closed simple subpolyhedron onto which W collapses and P is *locally flat* in W, that is, for each point $p \in P$ there exists a local chart (U, ϕ) of W around p such that $\phi(P \cap U)$ is contained in $\mathbb{R}^3 \subset \mathbb{R}^4$.

It follows from this definition that in the 3-dimensional slice, the pair $(\mathbb{R}^3 \cap \phi(U), \mathbb{R}^3 \cap \phi(U \cap P))$ is PL-homeomorphic to one of the models depicted in Figure 5.

For the sake of simplicity, from now on we will skip the PL prefix and all the homeomorphisms will be PL unless explicitly stated otherwise. Not every 4-manifold admits a shadow: a necessary and sufficient condition for W to admit one is that it has a handle decomposition containing no handles of index greater than 2; see [**4**, **37**]. This imposes restrictions on the topology of W. For instance, its boundary has to be a nonempty connected 3-manifold.

DEFINITION 3.3. A shadow of an oriented, closed 3-manifold M is a closed shadow P of a 4-manifold W with $M = \partial W$.

THEOREM 3.4 (Turaev [**37**]). *Any closed, oriented, connected 3-manifold has a shadow.*

EXAMPLE 3.5. The simple polyhedron $P = S^2$ is a shadow of $S^2 \times D^2$, and hence of the 3-manifold $S^2 \times S^1$. In this case, P is a surface whose self-intersection number in the ambient 4-manifold is zero. Consider now the disk bundle over S^2 with Euler number equal to 1, homeomorphic to a punctured $\mathbb{CP}^2$. The 0-section of the bundle is a shadow of the 4-manifold homeomorphic to P, and so P is a shadow of $\mathbb{CP}^2 - B^4$ and of its boundary: S^3.

The above example shows that the naked polyhedron by itself is not sufficient to encode the topology of the 4-manifold collapsing on it. Turaev described [**37**] how to equip a polyhedron embedded in a 4-manifold with combinatorial data called *gleams* that are sufficient to encode the topology of the regular neighborhood of the polyhedron in the manifold. A gleam is a

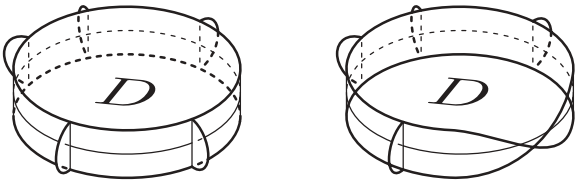

FIGURE 6. *An example of the set $U(D)$ in the case that D is a disc. Left: the band is an annulus. Right: the band is a Möbius strip; the combination is not embeddable in $\mathbb{R}^3$.*

coloring of the regions of the polyhedron with values in $\frac{1}{2}\mathbb{Z}$, with value modulo 1 given by a $\mathbb{Z}_2$-gleam which depends only on the polyhedron.

In the simplest case, if P is a shadow of M and P is homeomorphic to an orientable surface, then W is homeomorphic to an oriented disk bundle over the surface and the gleam of P is the Euler number of the normal bundle of P in W.

We summarize in the following proposition the basic construction of the $\mathbb{Z}_2$-gleam and of the gleam of a simple polyhedron. A *framing* for a graph G in a 3-manifold M is a surface with boundary embedded in M and collapsing onto G.

PROPOSITION 3.6. *Let P be a simple polyhedron. There exists a canonical $\mathbb{Z}_2$-coloring of the internal regions of P called the $\mathbb{Z}_2$-*gleam *of P. If P is embedded in a 4-manifold W in a flat way, there is a canonical coloring of the internal regions of P by integers or half integers called* gleams*, such that the gleam of a region of P is an integer if and only if its $\mathbb{Z}_2$-gleam is zero. Moreover, if $\partial P \subset \partial W$ is framed, then the gleam can also be defined on the noninternal regions of P.*

Proof. Let D be an internal region of P and let $\overline{D}$ be the abstract compactification of the (open) surface represented by D. The embedding of D in P extends to a map $i : \overline{D} \to P$ which is injective on $\text{int}(\overline{D})$, is locally injective on $\partial\overline{D}$, and sends $\partial\overline{D}$ into $\text{Sing}(P)$. In the case that i is injective, define a simple polyhedron $U(D)$ to be a small open regular neighborhood of $\overline{D}$ in P. In general, using i we can 'pull back' a regular neighborhood of $i(\partial\overline{D})$ and attach it to $\overline{D}$ to construct a simple polyhedron $U(D)$ that collapses onto $\overline{D}$. The map i extends to a map $i' : U(D) \to P$. The polyhedron $U(D)$ is constructed from $\overline{D}$ as follows:

(i) for each boundary component C of $\overline{D}$ take a band, either an annulus or a Möbius strip, and attach C to the core of the band; and

(ii) for each point $p \in \partial\overline{D}$ such that $i(p)$ is a vertex, take an arc A properly embedded in the corresponding band, meeting the core C transversally at p, and attach a half of the boundary of a disk to A.

An example in the case that D is a disk is shown in Figure 6. We define the $\mathbb{Z}_2$-gleam of D in P to be equal to the number of Möbius strip bands in the construction of $U(D)$, modulo 2. This coloring depends only on the combinatorial structure of P.

Let us now suppose that P is embedded in a 4-manifold W, and let D, $\overline{D}$, $i : \overline{D} \to P$, $U(D)$ and i' be defined as above. Using i', we can 'pull back' a neighborhood of $\overline{D}$ in W to an oriented 4-ball B^4 collapsing on $U(D)$. The regular neighborhood of a point $p_0 \in \partial\overline{D} \subset U(D)$ sits in a 3-dimensional slice B^3_0 of B^4 where it appears as in Figure 7. The direction along which the other regions touching $\partial\overline{D}$ get separated gives a section of the bundle of orthogonal directions to D in B^4. (If $p_0 \in \partial P$, use the framing of ∂P in place of the other regions. By an *orthogonal direction,* we mean a line in the normal bundle, not a ray.) This section can be defined on all $\partial\overline{D}$ and the obstruction to extending it to all of $\overline{D}$ is an element of $H^2(\overline{D}, \partial\overline{D}; \pi_1(S^1))$. Since B^4 is oriented, we can canonically identify this element with an integer and define the gleam of D to be half of this number. □

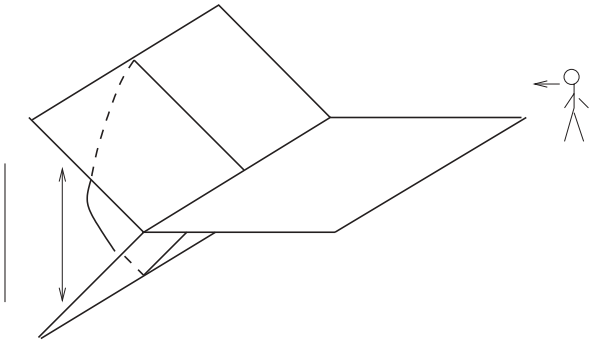

FIGURE 7. *The picture sketches the position of the polyhedron in a 3-dimensional slice of the ambient 4-manifold. The horizontal plane is our region of interest D. The direction indicated by the vertical double arrow is the one along which the two regions touching the horizontal one get separated.*

We can also go in the other direction, from gleams to 4-manifolds.

DEFINITION 3.7. A gleam on a simple polyhedron P is a coloring on all the regions of P with values in $\frac{1}{2}\mathbb{Z}$, such that the color of each internal region is an integer if and only if its $\mathbb{Z}_2$-gleam is zero.

THEOREM 3.8 (Reconstruction of a 4-manifold; Turaev [**37**]). *Let P be a polyhedron with gleams g; there exists a canonical reconstruction associating to P and g a pair (W_P, P). Here W_P is a PL, compact, oriented 4-manifold containing a properly embedded copy of P with framed boundary, so that W_P collapses onto P and so that the gleam of P in W_P coincides with g. The pair (W_P, P) can be explicitly reconstructed from P and its gleam. Moreover, if P is a polyhedron embedded in a PL, compact, oriented 4-manifold W, ∂P is framed, and g is the gleam induced on P as explained in the Proposition 3.6, then W_P is homeomorphic to the regular neighborhood of P in W.*

Hence, to study 4-manifolds with shadows (and their boundaries), one can either use abstract polyhedra equipped with gleams or embedded polyhedra.

From now on, each time we speak of a shadow of a 3-manifold as a polyhedron, we will be implicitly taking a 4-dimensional thickening of this polyhedron whose boundary is the given 3-manifold or, equivalently, a choice of gleams on the regions of the polyhedron.

EXAMPLE 3.9. Let M be a 3-manifold which collapses onto a simple polyhedron P whose regions are orientable surfaces; it is straightforward to check that the $\mathbb{Z}_2$-gleam of P is everywhere zero. Let us then equip P with the gleam which is zero on all the regions; Turaev's thickening construction produces the 4-manifold $W = M \times [-1, 1]$ and P is a shadow of ∂W, homeomorphic to the double of M.

3.2. *Shadows of framed graphs in manifolds with boundary*

In this section, we extend the definition of shadows to pairs (M, G) where M is a 3-manifold, possibly with boundary, and $G \subset M$ is a (possibly empty) framed graph with trivalent vertices and univalent ends. To do this, we allow the simple polyhedron to have boundary.

DEFINITION 3.10. A *boundary-decorated* simple polyhedron P is a simple polyhedron where ∂P is equipped with a cellularization whose 1-cells are colored with one of the following colors: i (*internal*), e (*external*), and f (*false*). We can correspondingly distinguish three subgraphs of ∂P intersecting only in 0-cells and whose union is ∂P: let us call them $\partial_i P$, $\partial_e P$, and $\partial_f P$. A boundary-decorated simple polyhedron is said to be *proper* if $\partial_f(P) = \emptyset$.

We can turn decorated polyhedra into shadows. The intuition is that $\partial_f(P)$ is ignored, $\partial_e(P)$ is drilled out to create the boundary of the 3-manifold, and $\partial_i(P)$ gives a trivalent graph.

DEFINITION 3.11. Let P be a boundary-decorated simple polyhedron, properly embedded in a 4-manifold W which collapses onto P with a framing on $\partial_i(P)$. Let M be the complement of an open regular neighborhood of $\partial_e P$ in ∂W, and let G be a framed graph whose core is $\partial_i(P)$. Then we say that P is a *shadow* of (M, G) and, if $\partial_f P = \emptyset$, we call it a *proper shadow*. As before, we can define a *gleam* on each region of P that does not meet $\partial_f(P) \cup \partial_e(P)$.

REMARK 3.12. Turaev's reconstruction theorem extends to the case of decorated polyhedra equipped with gleams on the regions not touching $\partial_e P \cup \partial_f P$.

REMARK 3.13. The genus of a boundary component of M equals the rank of H_1 of the corresponding component of $\partial_e(P)$, since the Euler characteristic of the handlebody filling the boundary component is equal to the Euler characteristic of the graph. In particular, components of $\partial_e(P)$ corresponding to sphere boundary components of M are contractible, and so if $\partial_f(P)$ is empty, M has no sphere boundary components that do not meet G.

THEOREM 3.14 (Turaev [37]). *Let M be an oriented, connected 3-manifold and let G be a properly embedded framed graph in M with vertices of valence 1 or 3. If M has no spherical boundary components that do not meet G, the pair (M, G) has a proper, simply connected shadow.*

Let us now show how to construct a shadow of a pair (M, G) given a shadow P of $(M, \emptyset)$. Recall that M is the boundary of a 4-manifold collapsing through a projection π onto P. Up to small isotopies, we can suppose that the restriction to G of π is transverse to $\mathrm{Sing}(P)$ and to itself; that is, it does not contain triple points or self-tangencies and is injective on the vertices of T. Let us also suppose that it misses $\partial_f(P)$. Then the mapping cylinder of the projection of G in P is contained in the thickening W_P of P and W_P collapses on it. (Recall that the mapping cylinder is $P \cup G \times [0, 1]$, with $G \times \{0\}$ identified with $\pi(G) \subset P$.) By Proposition 3.6, we can equip this polyhedron with gleams. Coloring $G \times \{1\}$ with the color i, we get a shadow of the pair (M, G), coloring it with e we get a shadow of $M - U(G)$, where $U(G)$ is a small open regular neighborhood of G in M, and coloring it with f we get another shadow of M (necessarily not proper).

As a warm up, note that a flat disk D whose boundary has color f is a shadow of the pair $(S^3, \emptyset)$. The open solid torus $T_h = \pi^{-1}(\mathrm{int}(D))$ can be imagined as the regular neighborhood of the closure of the z-axis in $\mathbb{R}^3$, embedded inside S^3 in the standard way. The fibers of $\pi : S^3 \to D$ run parallel to the z-axis away from ∞ and are unknotted. The projection of the solid torus $T_v = S^3 \backslash T_h$ is ∂D.

With the setup above, we now apply the projection construction to the case of a link L in S^3. Up to isotopy, we can suppose that $L \subset T_h$ and that its projection to D is generic; so it is sufficient to consider a standard diagram of L in the unit disk in $\mathbb{R}^2$. The mapping cylinder D_L of $\pi : L \to D$ is obtained from D by gluing an annulus for each component of L and marking the free boundary components of these annuli with the color i. We can further

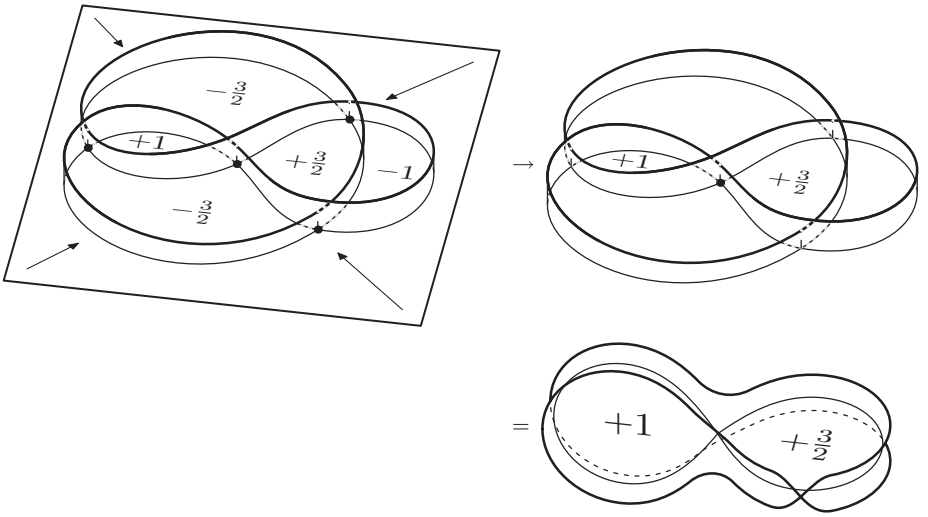

FIGURE 8. *In the left part of the picture, we sketch the construction described in Example 3.15; the resulting shadow is drawn in the right part where for the two internal regions we write their gleams. Note that, after the collapse of the polyhedron D_L along its free boundary component (as indicated by the arrows), the only vertex surviving is the central one.*

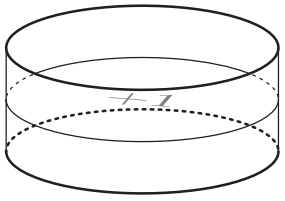

FIGURE 9. *The shadow obtained for the Hopf link.*

collapse the region of D_L containing $\partial_f D$; this produces a simple subpolyhedron of D_L, which we call P_L. By construction $\partial P_L = \partial_i P_L = L$.

In general, P_L has some vertices, each corresponding to a crossing in the diagram of L. However, some of the crossings in that diagram of L do not generate vertices in P_L because they disappear when we pass from D_L to P_L.

EXAMPLE 3.15. Applying the construction to a figure-8 knot in a standard position, one gets a shadow of its complement containing only one vertex: three of the four crossings of the diagram are contained in the boundary of the region to be collapsed in P_L. See Figure 8.

EXAMPLE 3.16. Consider a standard Hopf link in T_h. The polyhedron H that one gets by applying the above procedure contains no vertices: the two crossings of the standard projection of the Hopf link in $\mathbb{R}^2$ touch the region of D which is collapsed. The resulting polyhedron can be obtained by gluing a disk to the core of an annulus; its embedding in B^4 is such that ∂H is the Hopf link in ∂B^4. See Figure 9.

EXAMPLE 3.17. Applying this construction to the graph ∂P in Figure 20 (with the projection shown there) gives the shadow in Figure 21.

3.3. *Shadow complexity and its basic properties*

We will now define a notion of *shadow complexity* and study how it behaves under combining manifolds, either via connect sum (gluing along spheres) or torus connect sum (gluing along torus boundaries).

DEFINITION 3.18. For M to be an oriented 3-manifold (possibly with boundary) and $T \subset M$ a trivalent graph, the *shadow complexity* $\mathrm{sc}(M,T)$ of the pair (M,T) is the minimal number of vertices of a boundary-decorated shadow of (M,T).

REMARK 3.19. If M has no spherical boundary components, it does not matter whether or not we allow the shadow to have false edges in this definition: if we have a decorated shadow P_1 for (M,T), it can be shown that the polyhedron P_2 obtained by iteratively collapsing all the regions of P_1 containing a false boundary edge is a complex obtained by gluing some graphs to a (possibly disconnected) simple polyhedron. This complex P_2 can be modified, without adding vertices, to give a shadow P_3 for (M,T) without false edges and no more vertices than P_1; the modifications are generally similar to those in Lemma 3.22, with a few special constructions for cases where the complex is contractible (so M is S^3) or the graph has nontrivial loops, producing $S^1 \times S^2$ summands in the prime decomposition. All of these special cases are graph manifolds. By Proposition 3.31, they can be treated without creating any vertices.

Such a notion of complexity is similar to the usual notion of complexity of 3-manifolds introduced by S. Matveev [**27**].

DEFINITION 3.20. The *complexity* $c(M)$ of a 3-manifold M is the minimal number of vertices in a simple polyhedron P contained in M which is a spine for M or M minus a ball.

Both notions are based on the least number of vertices of a simple polyhedron describing (in a suitable sense) the given manifold. Despite this similarity, shadow complexity is not finite. That is, the set of manifolds having complexity less than or equal to any given integer is infinite. For instance, the lens spaces $L(p,1)$ have a shadow surface which is S^2 with gleam p, and so they all have shadow complexity 0.

To reduce the set of attainable manifolds to a finite number and bound the complexity of the 4-manifold, we also need to bound the gleams.

DEFINITION 3.21. The *gleam weight* $|g|$ of a shadow polyhedron (P,g) is the sum of the absolute values of the gleams on the regions of P.

LEMMA 3.22. *Shadow complexity is subadditive under connected sum: for M_1, M_2 two oriented 3-manifolds containing graphs T_1, T_2,*

$$\mathrm{sc}(M_1 \# M_2, T_1 \cup T_2) \leqslant \mathrm{sc}(M_1,T_1) + \mathrm{sc}(M_2,T_2).$$

Proof. Let P_1 and P_2 be two shadows for (M_1,T_1) and (M_2,T_2) having the least number of vertices, and let W_1 and W_2 be the corresponding 4-thickenings. To construct a shadow of the connect sum, let x_1 and x_2 be two points in regions of P_1 and P_2, respectively, and join them by an arc. The polyhedron that we get can be embedded as a shadow of the boundary connected sum of W_1 and W_2. This polyhedron is not simple, so we modify the construction slightly: roughly speaking, we put our fingers at the two ends of the arc and push P_1 toward P_2 along the arc until they meet in the middle along a disk. More precisely, we identify a closed regular neighborhood of x_1 and x_2, and put gleam 0 on the resulting disk region. □

QUESTION 3.23. Is shadow complexity additive under connected sum?

If the answer to the above question were 'yes', a consequence would be the following.

LEMMA 3.24. *If shadow complexity is additive under connected sum, then for any closed 3-manifold M,*

$$\mathrm{sc}(M) \leqslant 2c(M),$$

where $c(M)$ is Matveev's complexity.

Proof. Let P be a minimal spine of M, that is, a simple polyhedron whose 3-thickening is homeomorphic to the complement M' of a ball in M and containing the least possible number of vertices. Then P, equipped with gleam 0 on every region, is a shadow of $M' \times [-1, 1]$, with boundary $M \# \overline{M}$. Therefore, $\mathrm{sc}(M \# \overline{M}) \leqslant c(M)$ and the thesis follows. □

It is worth noting that the consequence of the above lemma is true for all the 3-manifolds with Matveev's complexity up to 9: we were able to check the inequality for all of them using Proposition 3.27 and the basic blocks exhibited by Martelli and Petronio [**26**].

We next show that shadow complexity does not increase under Dehn surgery.

LEMMA 3.25. *Let L be a framed link contained in an oriented 3-manifold M and let P be a shadow of (M, L). A manifold M' obtained by Dehn surgery on L has a shadow obtained by capping each component of ∂P by a disk.*

Proof. Let W be a 4-thickening of P. Surgery of M along a component of L with integer coefficients corresponds to gluing a 2-handle to W. Gluing the core of this 2-handle to P gives a shadow of W, and the definition of Dehn surgery on a framed link ensures that the gleam on the capped region does not change. □

REMARK 3.26. Lemma 3.25, together with the projection construction described in Section 3.2, gives an easy proof that any closed 3-manifold has a shadow, since any 3-manifold can be presented by an integer surgery on a link in S^3.

PROPOSITION 3.27. *Let M_1 and M_2 be two oriented manifolds such that both ∂M_1 and ∂M_2 contain torus components T_1 and T_2. Let P_1 and P_2 be shadows of M_1 and M_2, and let M be any 3-manifold obtained by identifying T_1 and T_2 with an orientation-presevering homeomorphism. Then M has a shadow which can be obtained from P_1 and P_2 without adding any new vertices. In particular, any Dehn filling of a 3-manifold can be described without adding new vertices.*

Proof. Let W_1 and W_2 be the 4-thickenings of P_1 and P_2. The tori T_1 and T_2 are equipped with the meridians μ_1 and μ_2 of the external boundary components l_1 and l_2 of P_1 and P_2. Also fix longitudes λ_i on them. The orientation-reversing homeomorphism identifying T_2 and T_1 sends μ_2 into a simple curve $a\lambda_1 + b\mu_1$ and λ_2 into a curve $c\lambda_1 + d\mu_1$.

We now describe how to modify P_1 and construct a shadow P_1' of M_1 embedded in a new 4-manifold W_1', such that the meridian induced by P_1' on T_1 is the curve $a\lambda_1 + b\mu_1$. To construct P_1', let us construct a shadow of the Dehn filling of M_1 along T_1 whose meridian is $a\lambda_1 + b\mu_1$. It is a standard fact that any surgery on a framed knot can be translated into an integer surgery over a link as shown in Figure 10.

With the notation of the figure, we glue n copies of the polyhedron H of Example 3.16 to P_1, so that one component of ∂H_1 is identified with l_1, H_j is glued to H_{j+1} and H_{j-1}, and the free

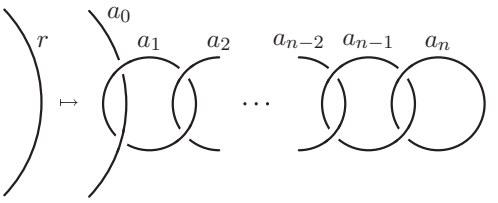

FIGURE 10. *In this picture we show how to transform a rational surgery with coefficient r over a knot into an integer surgery over a link. The a_i are the coefficients of a modified continued fraction expansion of r, in the form $r = a_0 - 1/(a_1 - (1/a_2 \cdots - 1/a_n))$.*

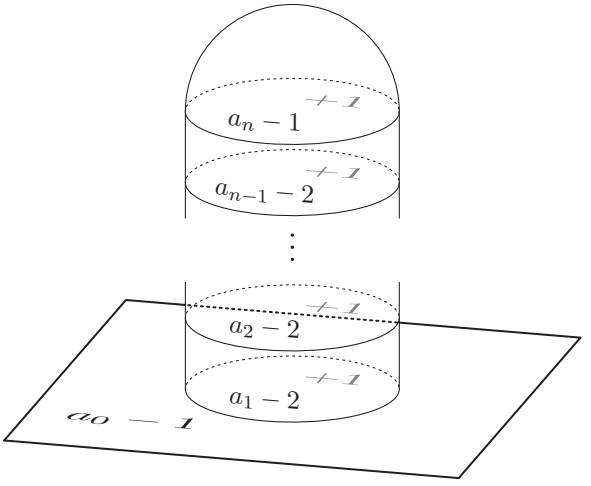

FIGURE 11. *The shadow with no vertices corresponding to the surgeries on the chain in Figure 10. Intrinsically in the 4-manifold, this is equivalent to a chain of spheres, each intersecting the next in just one point.*

component of ∂H_n is a knot l_1'. On the level of the boundary of the thickening, we are gluing n copies of the complement of the Hopf link in S^3 to M. Since the complement of the Hopf link is $T^2 \times [0,1]$, the final 3-manifold is unchanged. But now the polyhedron $P_1' = P_1 \cup H_1 \cup \ldots \cup H_n$ can be equipped with gleams so to describe the operation of Figure 10. The meridian μ_1' of l_1' is now by construction the curve which, expressed in the initial base of T_1, is $a\lambda_1 + b\mu_1$. Hence, we can now glue P_1' and P_2 along l_1' and l_2 and choose suitably the gleam of the region of $P_1' \cup P_2$ to obtain the desired homeomorphism. □

COROLLARY 3.28. *Shadow complexity is subadditive under torus sums, that is, under the gluing along toric boundary components through orientation-reversing homeomorphisms.*

One special case of torus sums is surgery, gluing in a solid torus. Surgery can decrease any reasonable notion of complexity, so shadow complexity cannot be additive under torus sums in general; hence we ask the following question.

QUESTION 3.29. Let M_1 and M_2 be two oriented 3-manifolds with incompressible torus boundary components T_1 and T_2. Is it true that any torus sum of M_1 and M_2 along T_1 and T_2 has shadow complexity equal to $\mathrm{sc}(M_1) + \mathrm{sc}(M_2)$?

3.4. *Complexity zero shadows*

In this section, we classify the manifolds having zero shadow complexity.

DEFINITION 3.30. An oriented 3-manifold is said to be a *graph manifold* if it can be decomposed by cutting along tori into blocks homeomorphic to solid tori and $R \times S^1$, where R is a pair of pants (that is, a thrice-punctured sphere).

Graph manifolds can also be characterized as those manifolds which have only Seifert fibered or torus bundle pieces in their JSJ decomposition.

PROPOSITION 3.31 (Complexity zero manifolds). *The set of oriented* 3*-manifolds admitting a shadow containing no vertices coincides with the set of oriented graph manifolds.*

Proof. To see that any graph manifold has a shadow without vertices, notice that a disk with boundary colored by e is a shadow of a solid torus and, similarly, a pair of pants R is a shadow of $R \times S^1$. Proposition 3.27 shows that any gluing of these blocks can be described by a shadow without vertices.

For the other direction, we must show that if a 3-manifold M has a shadow P without vertices, then it is a graph manifold. The polyhedron P can be decomposed into basic blocks as follows. Since P contains no vertices, a regular neighborhood of $\mathrm{Sing}(P)$ in P is a disjoint union of blocks of the following three types:

(i) the product of a Y-shaped graph and S^1;
(ii) the polyhedron obtained by gluing one boundary component of an annulus to the core of a Möbius strip; and
(iii) the polyhedron obtained by considering the product of a Y-shaped graph and $[-1, 1]$ and identifying the graphs $Y \times \{1\}$ and $Y \times \{-1\}$ by a map which rotates the legs of the graph of $\frac{2\pi}{3}$.

Let $\pi : M \to P$ be the projection of M on P. The complement of the above blocks in P is a disjoint union of (possibly nonorientable) compact surfaces. The preimage under π of each of these surfaces is a (possibly twisted) product of the surface with S^1, and hence is a graph manifold. Moreover, the preimage under π of the above three blocks is a 3-dimensional submanifold of M which admits a Seifert fibration (induced by the direction parallel to $\mathrm{Sing}(P)$), and hence is a graph manifold. □

3.5. *Decomposing shadows*

In Proposition 3.31, we saw how to decompose a shadow with no vertices into elementary pieces. For more general shadows, we will need a new type of block. For simplicity, we will suppose that the boundary of P is all marked 'external' and that the singular set $\mathrm{Sing}(P)$ of the shadow P is connected and contains at least one vertex. (The latter statement can always be achieved by modifying P with suitable local moves.) Let P be a shadow for a 3-manifold M, possibly with nonempty boundary, and let $\pi : M \to P$ be the projection. Then we have the following.

PROPOSITION 3.32. *The combinatorial structure of* P *induces through* π^{-1} *a decomposition of* M *into blocks of the following three types:*

(i) *products* $F \times S^1$*, where* F *is an orientable surface, or* $F \tilde{\times} S^1$ *with* F *nonorientable;*
(ii) *products of the form* $R \times [-1, 1]$*, where* R *is a pair of pants; and*
(iii) *genus-*3 *handlebodies.*

Proof. Decompose the polyhedron P by taking regular neighborhoods of the vertices and then regular neighborhoods of the edges in the complement of the vertices. This decomposes P into blocks of the following three types:

(i) surfaces (corresponding to the regions);
(ii) pieces homeomorphic to the product of a Y-shaped graph and $[-1,1]$; and
(iii) regular neighborhoods of the vertices.

The preimage of the first of these blocks is a block of the first type in the statement. Let us consider the preimages of the products $Y \times [-1,1]$. The 4-dimensional thickening of one of these blocks is the product of the 3-dimensional thickening $\mathbf{Y}$ of the Y-graph and $[-1,1]$, where $\mathbf{Y}$ is a 3-ball containing a properly embedded copy of Y and collapsing on it. The preimage in M of this block is the product of $[-1,1]$ with $\partial\mathbf{Y} - \partial Y$, which is a pair of pants.

Let us denote by V the simple polyhedron formed by a regular neighborhood of a vertex in P. We are left to show that $\pi^{-1}(V)$ is a genus-3 handlebody. The 4-thickening of V is $\mathbf{V} \times [-1,1]$, where $\mathbf{V}$ is the 3-dimensional thickening of V, that is, a 3-ball into which V is properly embedded. In particular, $\partial V \subset \partial\mathbf{V}$ is a tetrahedral graph and so $\partial\mathbf{V}$ is split into four disks by ∂V. One can decompose $\partial(\mathbf{V} \times [-1,1])$ as $\partial\mathbf{V} \times [-1,1] \cup \partial\mathbf{V} \times \{-1,1\}$. The part of this boundary corresponding to M is the complement of ∂V and is homeomorphic to $\mathbf{V} \times \{-1\} \cup (\partial\mathbf{V} - \partial V) \times [-1,1] \cup \mathbf{V} \times \{1\}$; this is composed of two 3-balls connected through four handles of index 1 (each of which corresponds to one of the four disks into which ∂V splits $\partial\mathbf{V}$). □

Note that the boundaries of the blocks of the second two types in Proposition 3.32 are themselves naturally decomposed into annuli and pairs of pants.

3.6. *A family of universal links*

Now suppose further that P (and therefore M) has no boundary, and consider the union of the blocks of the second two types in Proposition 3.32. These two types of blocks meet in pairs of pants, and the remaining boundary is obtained from the annuli; therefore, we are left with a manifold S_P with boundary a union of tori, which depends only on the polyhedron P and not on the gleams. (The original manifold M can be obtained by surgery on S_P.)

In this section, we show that S_P is a hyperbolic cusped 3-manifold whose geometrical structure can be easily deduced from the combinatorics of P. We furthermore show how to present S_P as the complement of a link in a connected sum of copies of $S^2 \times S^1$.

As before, let P be a simple polyhedron (now with no boundary) such that $\mathrm{Sing}(P)$ is connected and contains at least one vertex; let $c(P)$ be the number of vertices. Let $S(P)$ be the regular neighborhood of $\mathrm{Sing}(P)$ in P, which we think of as a simple polyhedron with boundary colored 'internal'. Let $l_1, \ldots, l_k$ be the components of $\partial S(P)$ in P. To each l_i we assign a positive integer number c_i called its *valence* by counting the number of vertices touched by the region R_i of $S(P)$ containing l_i and an element of $\mathbb{Z}_2$ given by the $\mathbb{Z}_2$-gleam g_i of the region of $S(P)$ containing l_i.

Let X_P be the 4-thickening of $S(P)$ provided by Turaev's reconstruction theorem; X_P collapses onto a graph with Euler characteristic $\chi(S(P)) = -c(P)$, and so ∂X_P is a connected sum of $c(P)+1$ copies of $S^2 \times S^1$. Moreover, $\partial S(P)$ is a link L_P in ∂X_P. The manifold S_P introduced earlier is the complement of L_P in ∂X_P.

S_P has a natural hyperbolic structure which we can understand in detail, as we will now see.

PROPOSITION 3.33. *For any standard shadow surface P, S_P can be equipped with a complete, hyperbolic metric with volume equal to $2v_{\mathrm{oct}}c$.*

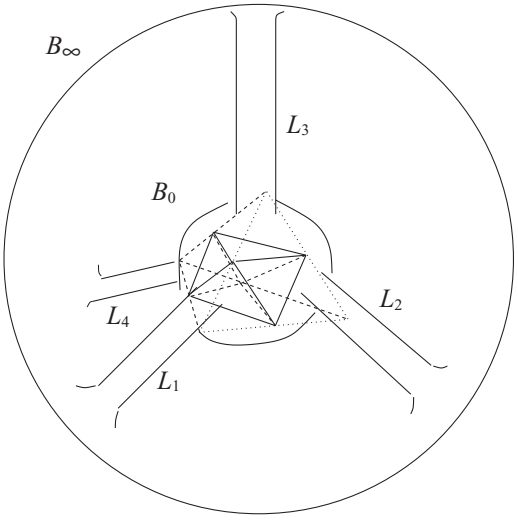

FIGURE 12. *In this picture we show how to connect the two balls B_0 and B_∞ in S^3 using the four legs L_i, $i = 1, \dots, 4$. In the center of B_0, we visualize how the regular octahedron O_0 is embedded. In the figure, the four internal faces are directed toward the four legs of the handlebody, since they are identified with the four internal faces of the octahedron O_∞.*

Proof. The main point of the proof is to construct a hyperbolic structure on a block corresponding to a vertex in $S(P)$ and then to show that these blocks can be glued by isometries along the edges of $S(P)$.

Let us realize a block of type 3 as follows. In S^3, pick two disjoint 3-balls B_0 and B_∞ forming neighborhoods, respectively, of 0 and ∞. Connect them using four 1-handles L_i, $i = 1, \dots, 4$, positioned symmetrically, as shown in Figure 12.

In the boundary of the genus-3 handlebody so obtained, consider the four thrice-punctured spheres formed by regular neighborhoods of the theta-curves connecting B_0 and B_∞, each of which is formed by 3-segments parallel to the cores of three of the 1-handles. These four pants are the surfaces onto which the blocks of type 2 in Proposition 3.32 are to be glued. Indeed, these blocks are of the form $R \times [-1, 1]$ where R is a thrice-punctured sphere, and they are glued to the blocks of type 3 along $R \times \{-1, 1\}$. We will now exhibit a hyperbolic structure on this block, so that these four thrice-punctured spheres become totally geodesic and their complement is formed by six annuli which are cusps of the structure.

Consider a regular tetrahedron in B_0 whose barycenter is the center of B_0 and whose vertices are directed in the four directions of the 1-handles L_i. Truncate this tetrahedron at its midpoints as shown in Figure 12. The result is a regular octahedron O_0 contained in B_0, with four faces (called 'internal') corresponding to the vertices of the initial tetrahedron, and four faces (called 'external') corresponding to the faces of the initial tetrahedron. Do the same construction around ∞ and call the result O_∞. The handlebody $B_0 \cup B_\infty \cup L_i$, $i = 1, \dots, 4$, can be obtained by gluing the internal faces of O_0 to the corresponding internal faces of O_∞. The remaining parts of the boundaries of the two octahedra are four spheres, each with three ideal points and triangulated by two triangles. If we put the hyperbolic structure of the regular ideal octahedron on both O_0 and O_∞, then after truncating with horospheres near the vertices, we get the hyperbolic structure we were searching for: the geodesic thrice-punctured spheres come from the boundary spheres without their cone points and the annuli are the cusps of the structure. Each (annular) cusp has an aspect ratio of $\frac{1}{2}$, since it is the union of two squares, the sections of the cusps of an ideal octahedron near a vertex. To show that these blocks can