manuscripta math. 88, 291 - 310 (1995)

manuscripta mathematica

A finite graphic calculus for 3-manifolds

Riccardo Benedetti – Carlo Petronio

Received February 13, 1995

In this paper we provide a presentation for compact oriented 3-manifolds with non-empty boundary up to orientation-preserving homeomorphism via a calculus on suitable finite planar graphs with extra structure (*decorated graphs*). Closed manifolds are included in this representation by removing a 3-ball.

Decorated graphs have an intrinsic geometric counterpart, as they are actually obtained by considering standard spines of the manifold and extra structure on them (*decorated spines*). The calculus on graphs is derived from the Matveev-Piergallini moves on standard spines ([2], [6], [7], [9]) which we re-examine and adapt to our setting (in particular in Section 1 we establish an oriented theory).

A comparison with the presentation of closed 3-manifolds via surgery on framed links in S^3 and the corresponding Kirby calculus [3], [4] allows to single out peculiarities of our graphic calculus. If one represents links by generic projections there are formal analogies between the two presentations, in particular both are supported by quadrivalent planar graphs with simple normal crossings and both express the change of orientation on manifolds by a simple involution on the set of graphs. The main differences are that edge-colours are taken in $\mathbb{Z}$ for framed links and in $\mathbb{Z}_3$ for us, and that our calculus is generated by a *finite number* (in a strict sense) of *local moves*, while on one hand the band move of Kirby calculus is not local, and on the other hand the local general Kirby move depends on the arbitrarily large number of strands involved.

The finiteness of our calculus can be exploited to construct *polynomial invariants* of spines and 3-manifolds in a way formally very close to the elementary definition of the Kauffman bracket invariant of framed links. For every choice of initial data (for which there is a wide freedom) the construction produces an ideal in a polynomial ring, explicitly given by a finite set of integral generators, and a process which to every decorated graph associates a polynomial. This polynomial is defined as a state sum which satisfies certain linear skein relations, and the class of the polynomial modulo the ideal is invariant under the calculus. This construction is widely discussed in [8], and non-triviality of the invariants produced is supported by the proof that Turaev-Viro invariants [11] appear in this framework, with a very *simple* choice of initial data.

Our calculus has also been used in [1] for a formal *algebraic* treatment of Roberts' approach to the Turaev-Walker theorem.

For the reader's convenience we state the main result which we will establish. We confine ourselves here to the case of oriented manifolds, because the presentation is particularly easy to describe in this case, but a similar graphic presentation is provided below also for non-oriented manifolds.

We start by introducing a few definitions. Consider a finite planar quadrivalent graph Γ with simple normal crossings, and assume that some vertices of Γ are marked. In the sequel given such a graph we will always call *edges* of Γ those obtained by ignoring the vertices which are not marked (in other words, edges are locally embedded segments with marked endpoints). Remark that edges might well not cover the graph.

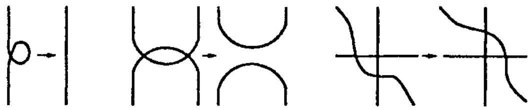

Figure 1: These moves are called respectively R_I, R_{II} and R_{III}

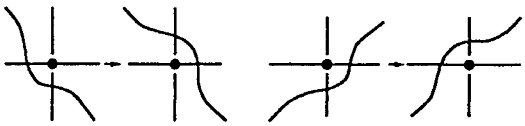

Figure 2: These moves are called respectively R'_{III} and R''_{III}

We define an o-graph to be the datum of:

1. A finite planar quadrivalent graph Γ with some marked vertices, where Γ is edge-connected in the sense just stated (in particular there is at least one marked vertex, and the edges cover Γ);
2. An under-over specification (as in the usual projections of links) at each marked vertex of Γ;
3. A colour chosen in $\mathbb{Z}_3$ attached to every edge of Γ.

We will always refer to Γ itself as an o-graph. We will establish the following:

Theorem 0.1. *To every o-graph there corresponds an oriented compact connected three-dimensional manifold with non-empty boundary well-defined up to orientation-preserving homeomorphism. Moreover:*

1. *Every such manifold is obtained from some o-graph with at least two marked vertices.*
2. *There exists an orientation-preserving homeomorphism between the manifolds associated to two o-graphs with at least two vertices if and only if the two o-graphs, regarded up to isotopies of the plane, are obtained from each other by a finite sequence of moves as shown in Figg. 1 to 4, and inverses of these moves.*
3. *Let an oriented manifold be associated to an o-graph. The same manifold with the opposite orientation is associated to an o-graph obtained as follows from the original one:*
   a. *At every marked vertex reverse the under-over specification.*
   b. *Change the colour of every edge to its inverse (in $\mathbb{Z}_3$).*

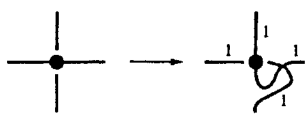

Figure 3: This move is called C. Here the convention is that to a certain edge it is equivalent to attach a colour i or many colours whose sum modulo 3 is i; in other words multiple colours can be summed up, and individual colours can be split into sums

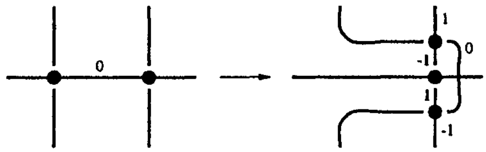

Figure 4: This move is called MP. Same convention as above about sums of colours

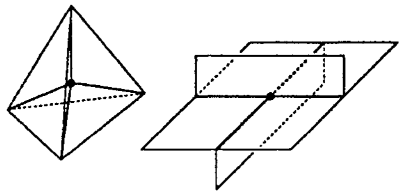

Figure 5: Two descriptions of the neighbourhood of a vertex in a standard spine

1 Standard spines and orientation

In this section we state a few facts about standard polyhedra and spines of 3-manifolds. Some of these facts were known, and we recall them with minor differences. On the other hand we could not find in the literature the oriented version of this theory: we are going to establish it here.

Throughout this work we operate in the PL category, without further explicit mention: the basic notions, including that of collapse which we will need, can be found in [10]. Manifolds will always be connected. A compact connected 2-dimensional polyhedron P is called quasi-standard if each point in it has a neighbourhood homeomorphic either to a plane, or to the union of three half-planes with common boundary line, or to the infinite cone over the 1-skeleton of a tetrahedron with vertex the barycentre of the tetrahedron (for further reference we show in Fig. 5 such a set). The singular set $S(P)$ and the set of vertices $V(P)$ of a quasi-standard polyhedron are naturally defined. We say that P is standard if the components of $P \setminus S(P)$ are discs and the components of $S(P) \setminus V(P)$ are segments.

We call P a standard spine of a compact 3-manifold M if P is a standard polyhedron and there exists an embedding $i : P \to \text{int}(M)$ such that M collapses onto $i(P)$. The following was established by Casler in [2] (see also [5] for a generalization to all dimensions):

Theorem 1.1. *Every compact 3-manifold with non-empty boundary admits a standard spine. Two 3-manifolds with homeomorphic standard spines are homeomorphic.*

According to this result if P is a standard spine of a manifold M we are allowed to define $M = \mathcal{M}(P)$. However $\mathcal{M}$ is not defined for all standard polyhedra (*cf.* Theorem 1.5 below). The following result was independently established by Matveev [6], [7] and Piegallini [9].

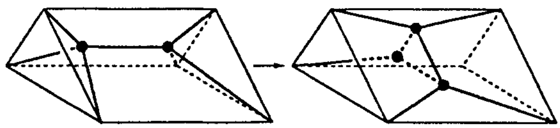

Figure 6: The Matveev-Piergallini move

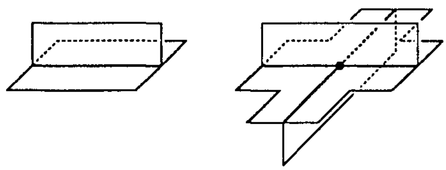

Figure 7: Two "typical shapes" occurring in a standard polyhedron (the third shape is the disc)

Theorem 1.2. *Every compact 3-manifold with non-empty boundary has a standard spine with at least two vertices. Standard spines with at least two vertices of homeomorphic manifolds are obtained from each other by a finite sequence of moves as shown in Fig. 6, and inverses of it.*

We want to characterize standard polyhedra which are spines. Before stating the result (Theorem 1.5) we need a few preliminaries.

A standard polyhedron can be canonically split into pieces of given shape. Namely, consider the objects shown in Fig. 7 and call them respectively $\mathfrak{E}$ and $\mathfrak{V}$. Remark that $\mathfrak{E}$ terminates with two triods while $\mathfrak{V}$ terminates with four triods (a triod is formally defined as the cone over three points). The following fact is straight-forward:

Lemma 1.3. *A standard polyhedron with n vertices can be obtained in an essentially unique way from n copies of $\mathfrak{V}$, $2n$ copies of $\mathfrak{E}$ and some copies of the disc, by identifying each terminal triod of the $\mathfrak{E}$'s with a terminal triod of the $\mathfrak{V}$'s, and then gluing the discs along the resulting boundary circles.*

In the sequel will denote by $\mathfrak{T}$ the discrete space containing three points.

Lemma 1.4. *To every local embedding $j : S^1 \to S(P)$ we can associate a $\mathfrak{T}$-bundle on S^1 denoted by B_j as explained in Fig. 8.*

Proof of 1.4. We only need to remark that, thanks to the symmetries of a typical neighbourhood of a vertex, when j crosses a vertex the situation can always be realized as shown on the right-hand side of Fig. 8. [1.4]

The first part of the following result was established (in a different form) in [5] and quoted in the present form in [9]. We include the proof for completeness and because the machinery is used elsewhere. The second part is apparently new.

Theorem 1.5. *Let P be a standard polyhedron.*

(A) P is a standard spine of a manifold if and only if for each disc D of $P \setminus S(P)$, if $j : S^1 \to S(P)$ describes how D is attached to $S(P)$, then the $\mathfrak{T}$-bundle B_j is trivial.

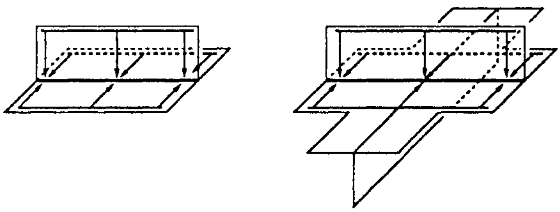

Figure 8: Definition of a $\mathfrak{T}$-bundle on S^1 associated to a local embedding j; the thick line is a portion of j and the little arrows define the projection of the bundle

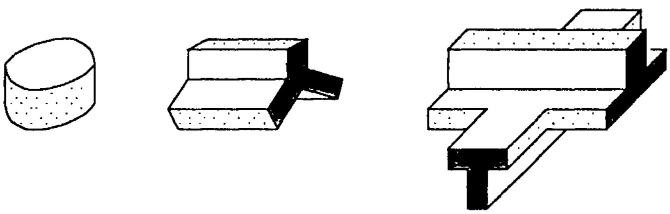

Figure 9: The thick pieces which glue up to give a neighbourhood of P

(B) P is the standard spine of an orientable manifold if and only if for every embedding $j : S^1 \to S(P)$ the $\mathfrak{T}$-bundle B_j has 1 or 3 components.

Proof of 1.5. (A) We start with the "only if" part. Let $\mathcal{J}$ be a finite union of circles and $\mathcal{T}$ be a finite union of triods such that $P \setminus (\mathcal{J} \cup \mathcal{T})$ is the splitting described in Lemma 1.3. Let D_i be the closure of the i-th disc component of $P \setminus (\mathcal{J} \cup \mathcal{T})$ $(i = 1, ..., d)$; D_i is a closed disc embedded in int(M); let W_i be a neighbourhood of D_i in M homeomorphic to a ball. If C_k is a non-disc component of $P \setminus (\mathcal{J} \cup \mathcal{T})$ let V_k be a neighbourhood of it in M also homeomorphic to a ball. We can arrange things in such a way that $W_i \cap V_k$ and $V_{k_1} \cap V_{k_2}$ are either empty or a ball, according as $D_i \cap C_k$ and $C_{k_1} \cap C_{k_2}$ are empty or not, while $W_{i_1} \cap W_{i_2}$ is always empty.

Now, operating within these balls, we can consistently thicken the D_i's and the C_k's, to obtain solid pieces as shown in Fig. 9. By "consistently" we mean the following: remark first that the boundary of each solid piece in Fig. 9 is divided into portions which can be white, black or dotted, and a dotted portion can be either an annulus or a strip; then "consistently" means that the black portions are identified in pairs, every dotted strip is glued to two other dotted strips along the terminal segments, the strips glued together give annuli, and each of these annuli is identified to some dotted annulus (the lateral surface of some cylinder).

Now fix i and let D_i be contained in the component D of $P \setminus S(P)$. We can say what it means for the $\mathfrak{T}$-bundle induced by the attaching function of D to be trivial: we abstractly glue the C_k's to ∂D_i, with each C_k appearing as many times as the number of components of $\mathcal{J} \cap C_k$, taking into account the identifications along $\mathcal{T}$, as in Fig. 10. (Remark that $\mathcal{J} \cap C_k$ can have up to 6 components if C_k is homeomorphic to $\mathfrak{Y}$, up to 3 otherwise.) Then triviality is the condition that the thick lines close up to two circles (three including the boundary of the disc).

Let us return to the situation of a neighbourhood of P in M covered by nicely

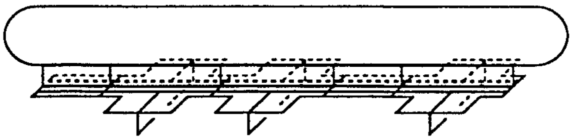

Figure 10: Pieces glued around a disc

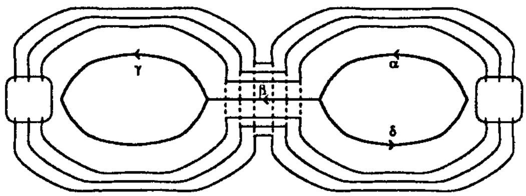

Figure 11: Within the box on the left the strands are joined in such a way that the bundle on γ has 1 or 3 components, and the global bundle has 2 components, hence within the box on the right the strands must be joined so that the bundle on $\alpha\delta$ has 2 components

glued solid pieces as shown in Fig. 9. We concentrate on the gluings involving the fixed disc D_i. If we give each of the pieces an arbitrary orientation and consequently orient the black portions of their boundary (see Fig. 9) the following conditions are pairwise equivalent:

1. The dotted strips (Fig. 9) glue up to a Möbius strip;
2. An odd number of identifications between black portions preserve the orientation;
3. The $\mathfrak{T}$-bundle is non-trivial.

Equivalences 1⇔2 and 2⇔3 are clear, and Not(1)⇒Not(3) implies the proof of "only if" in part A.

To establish "if" in part A, define a manifold M by gluing abstract solid pieces as in Fig. 9 as prescribed by the gluings in P, which can be done because of implication Not(3)⇒Not(1) (the pieces are now abstract and not in M, but the setting is the same). The result is a manifold which has standard spine P.

(B) Let P have the property that $\mathfrak{T}$-bundles induced on embeddings have either 1 or 3 components; we claim that the same property holds also for local embeddings. To see this, start with a local embedding j such that B_j has 2 components, and decompose j as $\alpha\beta\gamma\beta^{-1}\delta$, where γ is constant or an embedding of S^1 and α, β, δ are constant or local embeddings of $[0,1]$ in $S(P)$ (see Fig. 11). The same figure also proves that if one cuts out $\beta\gamma\beta^{-1}$ the bundle still has 2 components, and the conclusion follows because one eventually reaches a circle.

Our claim is proved. We deduce from it that P is the spine of a manifold: the $\mathfrak{T}$-bundle induced by the attaching function of a disc has a global section (a circle contained in the disc), so it cannot have 1 component; then it has 3 components, i.e. it is trivial.

So we can prove equivalence in part B assuming in any case that P is the standard spine of a 3-manifold M. We think of M as being built up from pieces as in Fig. 9. We recall that M is orientable if and only if there exists no loop $j : [0,1] \to M$ such that if the orientation is fixed at $j(0)$ and consistently followed along j, the orientation

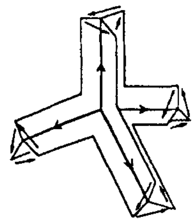

Figure 12: At every vertex the orientations along the four edges must be in this situation

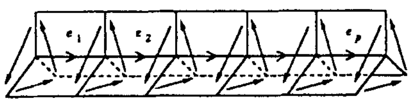

Figure 13: A loop cut open at a vertex

carried by $j(1)$ does not match the original one. Of course it is equivalent to confine oneself to embeddings of S^1 in $S(P)$. If one gives each non-cylindrical solid piece an arbitrary orientation and consequently orients the black portions of Fig. 9, the following statements for an embedding $j : S^1 \to S(P)$ are pairwise equivalent:

1. The loop j reverses the orientation (in the above sense);
2. An odd number of the identifications between black portions met along j preserve the orientation;
3. The $\mathfrak{T}$-bundle B_j over S^1 has two components.

Equivalences 1⇔2 and 2⇔3 are obvious, while 1⇔3 implies the conclusion. 1.5

We give now an intrinsic notion of orientation for a standard polyhedron. Fix P and $S = S(P)$. Remark that incident to every edge of S there are three germs of disc. We define an orientation of P along the edge as the choice of a direction for the edge and of a cyclic order for the three germs of discs, where a simultaneous reversal of the edge-direction and the disc-order is supposed to define the same orientation. At every edge there are exactly two orientations.

We define an orientation of P as the choice of an orientation of P along all the edges such that the compatibility condition shown in Fig. 12 holds at every vertex. Of course P admits either no orientation at all or exactly two.

Lemma 1.6. *A standard polyhedron P admits an orientation if and only if the $\mathfrak{T}$-bundle induced on every embedding of S^1 in $S(P)$ has either* 1 *or* 3 *components.*

Proof of 1.6. It is sufficient to show that if $e_1, ..., e_p$ are edges whose union is a circle then there exists a consistent orientation along them if and only if the $\mathfrak{T}$-bundle on the circle does not have 2 components. Direct the e_i's consecutively. If v is the first vertex of e_1 then of course the bundle is trivial on $e_1 \cup ... \cup e_p \setminus \{v\}$, and we can orient P along the e_i's so that the compatibility holds at the vertices different from v. Looking at the definitions one has the situation of Fig. 13, and then it is easily seen that the identification between the terminal triods preserves the cyclic order if and only if the resulting $\mathfrak{T}$-bundle does not have 2 components, whence the conclusion. 1.6

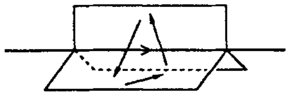

Figure 14: The rule to orient a spine

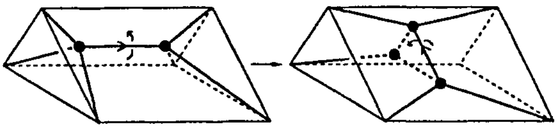

Figure 15: The oriented Matveev-Piergallini move

Corollary 1.7. *P is orientable if and only if it is the spine of an orientable manifold.*

This corollary is readily deduced from Lemma 1.6 and Theorem 1.5. But we have actually more: *there is a natural correspondence between orientations of the spine and orientations of the manifold.* This is obtained as follows. Let M be a manifold with a fixed orientation ω (we will denote it by M^ω). Let P be a standard spine of M and p be a point of some edge e of $S = S(P)$ and take a neighbourhood U of p in M and an orientation-preserving homeomorphism $\phi : U \to \mathbb{R}^3$. Define the orientation of P along e using ϕ and the usual left-handed screw rule in $\mathbb{R}^3$, as in Fig. 14. If u is the unique orientation of P extending this one along e, we set $M^\omega = \mathcal{M}^\Omega(P^u)$. Of course one has $\mathcal{M}^\Omega(P^{-u}) = M^{-\omega}$. Remark that Fig. 14 also provides the inverse rule: given an orientation of P, embed a neighbourhood of some point so that the situation is as in the figure, then define this embedding to be orientation-preserving and extend the orientation to M.

We define the oriented Matveev-Piergallini move on the set of oriented standard polyhedra in the natural way shown in Fig. 15 (this is a natural choice as there are points whose neighbourhood the move does not alter, and we require that the orientation near them is preserved). The following result is an analogue of Theorem 1.2 for the case of oriented manifolds.

Theorem 1.8. *1. To every oriented standard polyhedron P^u there corresponds an oriented compact manifold with non-empty boundary $M^\omega = \mathcal{M}^\Omega(P^u)$.*

2. For every such manifold M^ω there exists an oriented standard polyhedron P^u with at least two vertices such that $M^\omega = \mathcal{M}^\Omega(P^u)$.

3. Let $M_i^{\omega_i} = \mathcal{M}^\Omega(P_i^{u_i})$, $i = 1, 2$ where both P_1 and P_2 have at least two vertices. Then there exists an orientation-preserving homeomorphism of $M_1^{\omega_1}$ onto $M_2^{\omega_2}$ if and only if $P_1^{u_1}$ and $P_2^{u_2}$ are obtained from each other by a finite sequence of oriented Matveev-Piergallini moves and inverses of it.

Proof of 1.8. The first two assertions are readily deduced from Theorem 1.2, Theorem 1.5(B), Corollary 1.7 and the above construction. We prove the third assertion.

Let $P_1^{u_1}$ be obtained from $P_2^{u_2}$ by only one oriented Matveev-Piergallini move. Then we know that $M_1^{\omega_1}$ and $M_2^{\omega_2}$ are homeomorphic, but maybe not as oriented manifolds. We think of P_1 and P_2 as embedded in the manifold, each inducing on the manifold an orientation as stated above. As we have remarked, in the oriented move there are points whose oriented neighbourhoods are unchanged, and then the two orientations on the manifold must be the same. The case of more than one move is easily settled.

Now let $M_1^{\omega_1} \cong M_2^{\omega_2}$ with orientation. Then $P_1^{u_1}$ and $P_2^{u_2}$ are obtained from each other by unoriented moves. We can think of these moves as taking place within the same oriented manifold, in such a way that for every move there is a portion of the spine which is unchanged under it. We give the intermediate spines the orientation defined by their embedding in the manifold. Then all the Matveev-Piergallini moves are oriented. 1.8

2 Decorated spines and a graphic calculus for 3-manifolds: general case

As shown in the left-hand side of Fig. 5 a vertex of a standard polyhedron has a neighbourhood with a totally symmetric description; to be precise such a neighbourhood has 24 symmetries (only 12 if there is an orientation and we want it to be preserved). However it is sometimes useful to consider the less symmetric description of the right-hand side of Fig. 5; for instance we have used it in Fig. 8 to define the $\mathfrak{T}$-bundle induced on a singular loop. We can view in abstract terms the loss of symmetry at a vertex as a decoration of the vertex. Roughly speaking a decoration at a vertex is the choice of how to represent a neighbourhood of it as in the right-hand side of Fig. 5. To specify this representation we must say which of the six discs is "over" (the opposite one will be "under") and how the adjacent four discs must be arranged on the horizontal plane.

Let us be more formal. We will denote by $\mathfrak{V}_0$ the subset of $\mathbb{R}^3$ shown on the right-hand side of Fig. 5 (with the fixed embedding in $\mathbb{R}^3$ suggested by the figure). Let P be a standard polyhedron and let v be a vertex of P. We define a *decoration* of P at v as a pair (D,ω), where D is one of the six germs of discs at v and ω is an orientation for the union π of the four germs of discs adjacent to D (π is an open disc).

Let P be decorated at v by (D,ω). Let ϕ be a homeomorphism of a neighbourhood of v in P onto $\mathfrak{V}_0$. We say that ϕ is compatible with the decoration if $\phi(D)$ is the upper vertical sheet of $\mathfrak{V}_0$ and π is mapped in an orientation-preserving way to the union of the four horizontal sheets. Of course every neighbourhood of a vertex contains a set homeomorphic to $\mathfrak{V}_0$, and compatibility with the decoration essentially determines the homeomorphism (so the formal definition of decoration corresponds to the purpose it was meant for).

Remark 2.1. At every vertex a standard polyhedron admits precisely 12 decorations.

We define a decoration of a standard polyhedron as the choice of a decoration at every vertex. By definition a decoration of a standard polyhedron allows us to "canonically" embed in $\mathbb{R}^3$ neighbourhoods of the vertices. We consider now the problem of extending these embeddings to a neighbourhood of the singular set, which naturally leads to introducing a certain decorated planar graph. This discussion is very informal, but below we will provide precise statements.

Let P be decorated. We start with embeddings of neighbourhoods of the vertices compatible with the decoration, pushed apart by horizontal translations. Next, we extend the embedding to the whole singular set of P, with the requirement that the projection on the horizontal plane is generic. (Of course there is some arbitrariness in such an extension —we will take this into account.) Now we extend the embedding

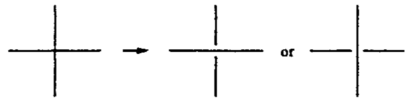

Figure 16: How to loosen a crossing which is not marked

across an edge: we will have triods arriving from both endpoints, and we only need to be careful and match them correctly in the middle. It might not be possible to do this in $\mathbb{R}^3$ (we just work with $\mathbb{R}^4$ in this case). Of course we do not have uniqueness (we could give either triod a complete twist before connection) but if we find an intrinsic method to number the three branches of each triod by $\{1,2,3\}$ we will have a well-defined permutation in $\mathfrak{S}_3$. Having found such a method our embedding will be essentially represented by a planar quadrivalent graph where some vertices are "fake", the others have some special decoration which allows us to recover the embedding, and the edges have a colour in $\mathfrak{S}_3$.

We will now formalize this construction. Since a graph defines a unique decorated standard polyhedron but the graph is not unique, it is convenient to start with graphs, show how they define standard polyhedra and then discuss non-uniqueness. Since we will use groups of permutations, we recall that the symbol $(i_1\, i_2\, \cdots\, i_k)$ denotes the function which maps i_j to i_{j+1} and i_k to i_1. We multiply two such symbols as functions, in the obvious way. For instance $(12)(13) = (132)$.

Let Γ be a finite connected quadrivalent graph embedded in $\mathbb{R}^2$ with normal crossings. Recall that by graph we mean a finite 1-complex. Assume that some vertices of Γ are marked, and define the edges of Γ as stated in the introduction by pretending that unmarked vertices do not exist.

We define an s-graph to be the datum of:

1. A finite planar quadrivalent graph Γ with some marked crossings, where Γ is edge-connected (in the sense just stated);
2. An under-over specification (as in the usual projections of links) at each marked vertex of Γ;
3. A direction for every edge of Γ;
4. A colour chosen in $\mathfrak{S}_3$ attached to every edge of Γ.

We will actually refer to Γ itself as an s-graph.

From an s-graph Γ to a decorated standard polyhedron: formal procedure.

1. Embed $\mathbb{R}^2$ in $\mathbb{R}^3$ as $\mathbb{R}^2 \times \{0\}$, and for each vertex of Γ which is not marked choose an arbitrary way in $\mathbb{R}^3$ to loosen it (there are two such ways, see Fig. 16).
2. For each marked vertex v select a small neighbourhood where Γ is a cross, and fit on it a subset of $\mathbb{R}^3$ as described in Fig. 17. We also use a planar notation which on one hand is readily deduced from the s-graph and on the other hand easily implies the 3-dimensional situation.
3. For each edge we select a small segment which is far from all vertices and associate to it a subset of $\mathbb{R}^3$ as described in Fig. 18. (Remark that for an end of an edge it is well-defined whether it is "over" or "under", so the definition makes sense.)
4. The objects defined in 2 and 3 terminate with T's which can be straight or upside-down. The choices we have made imply that we can extend the T's across the edges excluding the little segments, as shown in Fig. 19, by allowing horizontal

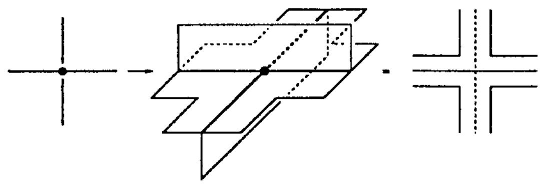

Figure 17: For every marked vertex (left) we introduce a 3-dimensional object (centre) and a planar symbol for it (right)

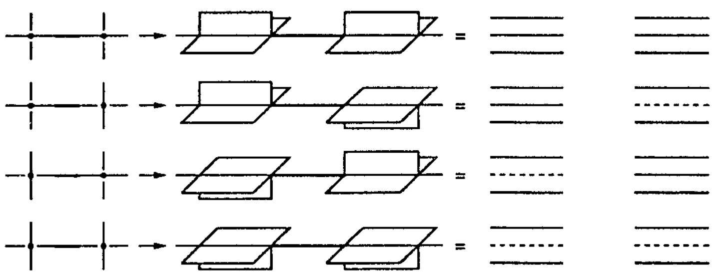

Figure 18: Depending on the four possible cases for the endpoints of the edge (left) we introduce a 3-dimensional object (centre) and a planar symbol for it (right)

rotations and vertical translations (these are necessary if we meet a loosened crossing, as in Fig. 20).

5. For each small segment we give a number 1, 2 or 3 to the three branches of each of the two T's involved as shown in Fig. 21. Remark that this numbering depends on the direction of the edge and on whether the T is straight or upside-down. We move the T across the little segment in such a way that i on the right is joined to $\sigma(i)$ on the left. It is quite easy to see that this can be done in $\mathbb{R}^3$ if and only if $\mathrm{sgn}(\sigma) = +1$; otherwise we just need to add one more dimension. Some examples are shown in Figg. 22 and 23.
6. We have obtained a polyhedron in $\mathbb{R}^3$ or at worst in $\mathbb{R}^4$ which has four types of points, three of them being the same as in a quasi-standard polyhedron. Points of the fourth type have neighbourhoods homeomorphic to half-planes but not to

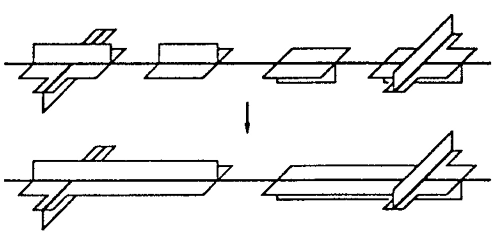

Figure 19: An example of how to extend the T's across an edge

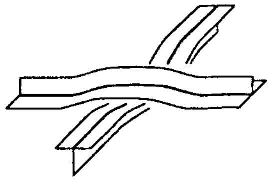

Figure 20: How to deal with a loosened vertex

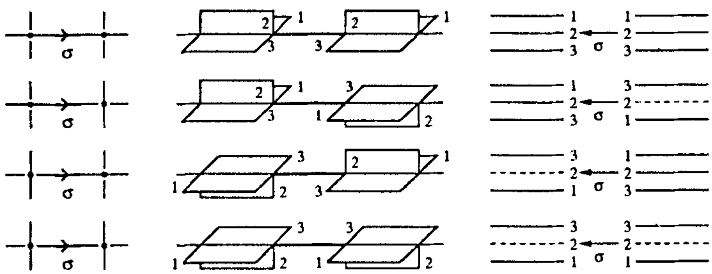

Figure 21: How to label the branches of the T's and a symbolic planar description of the gluings to be performed (step 5)

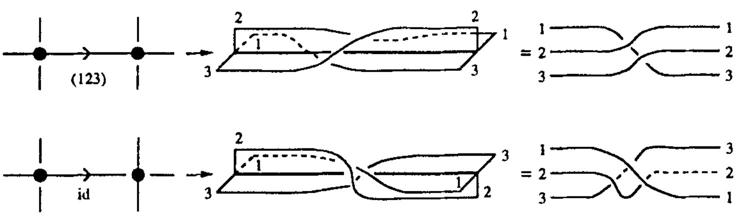

Figure 22: Two examples where everything takes place in the three-space

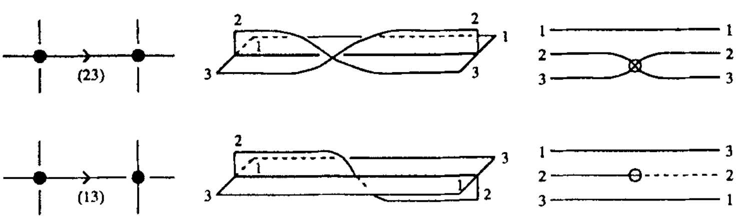

Figure 23: In these examples there are strips which cross each other without meeting, thanks to the fourth dimension not shown in the picture. A symbolic planar picture describing the situation is also provided

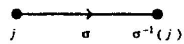

Figure 24: If i is one of $\{j,\sigma^{-1}(j)\}$ then i' is the other one, while $s'=\mathrm{sgn}(\sigma)\cdot s$

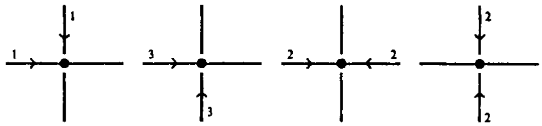

Figure 25: The two pairs (germ-of-edge,integer-near-it) give $\{(e,i),(e',i')\}$. The convention that when the direction of the edge is reversed the corresponding integer j must be replaced by $(1\,3)(j)$; it is easily checked that with this convention the four situations shown cover all possibilities

planes. The set of all points of the fourth type is a union of disjoint circles, and we attach a disc to each of them following the boundary.

7. We have obtained a standard polyhedron and an embedding in $\mathbb{R}^4$ of a neighbourhood of its singular part. The restriction of this embedding to a neighbourhood of a vertex, after an isotopy in $\mathbb{R}^3$ which preserves the horizontal plane, is a homeomorphism with the set E defined at the beginning of the section. Hence we can choose the decoration of the vertex with which this homeomorphism is compatible. We have therefore a decorated standard polyhedron.

Lemma 2.2. *If Γ is an s-graph the above-described procedure defines a unique decorated standard polyhedron $\mathcal{D}(\Gamma)$. Every decorated standard polyhedron arises like this.*

Proof of 2.2. Steps 2, 3 and 4 are unique up to isotopy in $\mathbb{R}^3$ (or $\mathbb{R}^4$); step 5 is not unique in $\mathbb{R}^3$ (or $\mathbb{R}^4$) because we can give complete twists, but the associated abstract object is unique; steps 6 and 7 are intrinsically defined; different choices in step 1 of course lead to the same abstract object. So the above construction defines a unique polyhedron, which is easily cheked to be standard.

The last assertion follows from the construction informally sketched above (start embedding near the vertices and extend). [2.2]

We will denote by $\mathcal{P}(\Gamma)$ the standard polyhedron obtained from $\mathcal{D}(\Gamma)$ by forgetting the decoration. Since we are interested in standard polyhedra as tools for studying 3-manifolds, we state the following:

Proposition 2.3. *Let Γ be an s-graph. Then $\mathcal{P}(\Gamma)$ is the standard spine of a manifold if and only if the following holds: Pick a germ of edge e and $i \in \{1,2,3\}$, set $s=+1$ and replace the triple $\{e,i,s\}$ by $\{e',i',s'\}$ alternatively as follows:*

1. *e' is the germ of the other end of the same edge and i', s' depend on the colour and direction of the edge, as described in Fig. 24.*
2. *e' is one of the three other germs at the same vertex, $s'=s$ and e', i' are given in the various possible cases by Fig. 25.*

Then the condition is that the first time we have $\{e,i,s'\}$ then we also have $s'=+1$, and this should happen for every initial choice of e and i.

Proof of 2.3. Using the procedure to recover $\mathcal{P}(\Gamma)$ from Γ one sees quite easily that this is exactly the condition that the $\mathfrak{T}$-bundle over the boundary of every disc is trivial (see Theorem 1.5). [2.3]

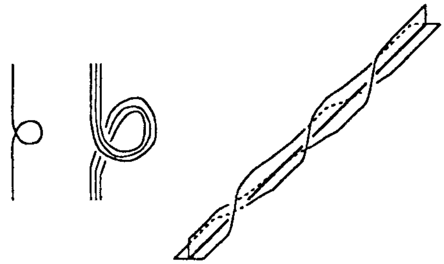

Figure 26: Proof that R_I represents a twist in the embedding

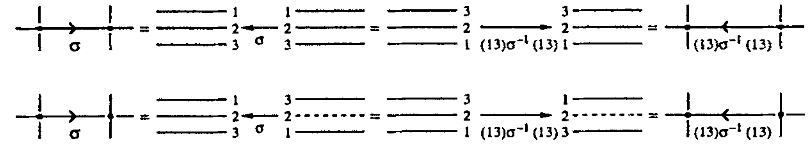

Figure 27: Reversing the direction: two examples

If $\mathcal{P}(\Gamma)$ is the standard spine of a manifold we call this manifold $\mathcal{M}(\Gamma)$.

We will discuss in the rest of this section the conditions under which two s-graphs define homeomorphic manifolds. There will be two intermediate steps corresponding to the questions: When do two s-graphs define the same decorated standard polyhedron? When do two s-graphs define the same standard polyhedron? The final result will be a set of moves with which one obtains from a given s-graph all and only the s-graphs defining the same manifold. In the next section we will deal with the case of oriented manifolds.

First of all we look at s-graphs up to homeomorphisms of the plane isotopic to the identity, without comments.

Proposition 2.4. *Given two s-graphs Γ_1 and Γ_2 there exists a decoration-preserving homeomorphism of $\mathcal{D}(\Gamma_1)$ onto $\mathcal{D}(\Gamma_2)$ if and only if Γ_1 and Γ_2 are obtained from each other by a finite sequence of the following moves:*

a. The "Reidemeister-type" moves R_I, R_{II}, R_{III}, R'_{III} and R''_{III} shown in Figg. 1 and 2. (We refer to a combination of them and their inverses as an R*-move.)*

b. The move U *which consists in reversing the direction of an edge and changing its colour from σ to $(13)\sigma^{-1}(13)$.*

Proof of 2.4. On considering the above construction, we can easily see that R-moves exactly recover arbitrariness for the embedding of a neighbourhood of the singular set (for an interpretation of R_I as a twist, see Fig. 26). We also had arbitrariness in choosing the directions of the singular edges, and we can show that U describes the natural way to colour an edge after reversing the direction, which implies the conclusion at once. We must distinguish the cases according as the endpoints of the edge are over or under. Figure 27 refers to the cases where both the endpoints are over and where the first endpoint is over and the second one is under. The other cases are treated similarly. 2.4

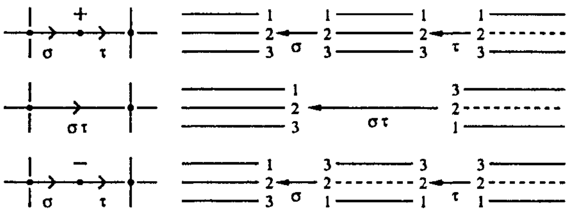

Figure 28: An example of product. This picture shows that whatever geometric interpretation one chooses for the subdivision point (top and bottom) the result is the product-move (centre)

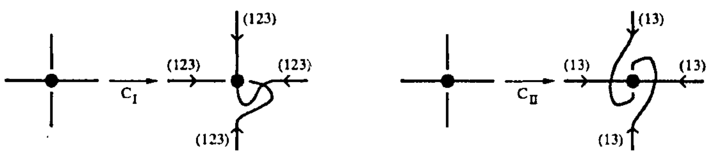

Figure 29: The moves C_I and C_{II}

Remark 2.5. The mapping $\mathfrak{S}_3 \ni \sigma \mapsto (13)\sigma^{-1}(13) \in \mathfrak{S}_3$ interchanges (12) and (23) and leaves the other elements unchanged.

From now on we will consider s-graphs up to moves R and U, i.e. we will view two s-graphs equivalent under them just as two different ways to draw the same object. With this convention an s-graph "is" a decorated standard polyhedron. Now, to discuss when $\mathcal{P}(\Gamma_1) = \mathcal{P}(\Gamma_2)$, we must see what is the effect on the s-graph of a change in the decoration of a vertex. This will result in the definition of new moves for s-graphs.

For the purpose of introducing these moves, we slightly extend the situation, by allowing the edges to be subdivided into finitely many subarcs, each with a direction and a colour. We introduce a product-move as follows: given two consecutive subarcs of an edge, we apply U (if necessary) to have them consecutively directed. Let the first subarc be decorated by σ and the second one by τ; the effect of the product-move is to remove the subdivision point between the arcs, direct the resulting arc in the obvious way and give it the colour $\sigma\tau$. To check that this definition is completely natural see the example of Fig. 28.

The moves we need, named C_I and C_{II} as they take place in the neighbourhood of a crossing, exploit the extension of the class of s-graphs which allows subdivisions. We think of the directions and colours of the edges involved as drawn outside the neighbourhood, so we do not deal with them. Both moves consist in adding four subdivision points near the vertex, joining these points to the vertex but not in the original way and then giving a certain direction and colour to the four arcs thus defined. The explicit description is given in Fig. 29; we will always tacitly assume that they must be followed by product-moves so that subdivisions disappear.

Remark that the move C_{II} is unambiguosly defined at every vertex, while there are two ways to apply C_I: in fact before the move the figure is symmetric under rotation through angle π, and after the move it is not. One could specify how to apply the move by stating which branch must be over both before and after the move.

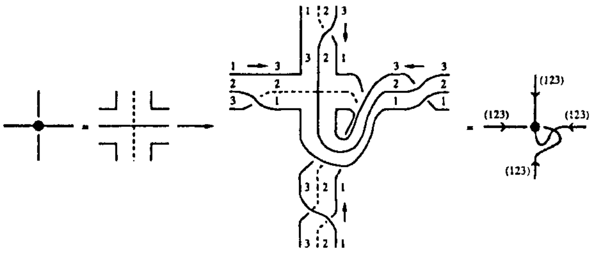

Figure 30: Explanation of the move C_I

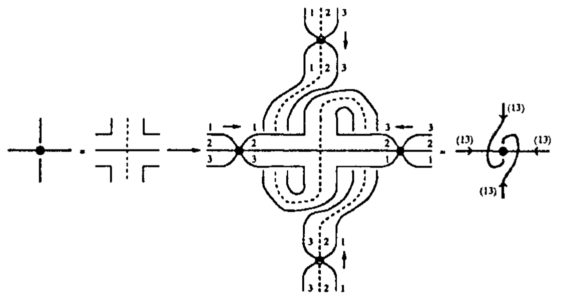

Figure 31: Explanation of the move C_{II}

Proposition 2.6. *If* Γ_1, Γ_2 *are s-graphs (regarded up to isotopy and moves* R *and* U*),* $\mathcal{P}(\Gamma_1)$ *and* $\mathcal{P}(\Gamma_2)$ *are homeomorphic if and only if* Γ_1 *and* Γ_2 *are obtained from each other by a finite sequence of moves* C_I *and* C_{II}.

Proof of 2.6. By Proposition 2.4 we must discuss how the decoration of a standard polyhedron can be changed. Figures 30 and 31 show that both the moves describe a modification of the embedding of a neighbourhood of a vertex, i.e. a modification of the decoration at the vertex.

For the conclusion of the proof it is sufficient to check that the two changes of decoration induced by C_I and C_{II} generate all 12 possible decorations at a vertex. To see this it is sufficient to describe the situation in a more intrinsic way. We recall that an abstract decoration is the choice of a germ of disc and an orientation for the union of the four discs which are adjacent to it. Then the effect of C_{II} on the decoration is of course just to reverse the orientation.

For C_I, fix a decoration (D_{12},ω) and denote by $e_1,...,e_4$ the germs of edges so that (e_1,e_3,e_2,e_4) are positively arranged with respect to ω, and D_{ij} is the germ of disc containing $e_i \cup e_j$ (see Fig. 32). Then the effect of the two ways to apply C_I is to replace (D_{12},ω) by (D_{24},ω') or (D_{13},ω''), where ω' extends $\omega|_{D_{14}}$ and ω'' extends $\omega|_{D_{23}}$. It is therefore evident that we can obtain all 12 decorations. $\boxed{2.6}$

Remark 2.7. By repeated composition of moves C_I and C_{II} at every vertex exactly 12 different situation can be produced. This is obvious if one thinks of the intrinsic

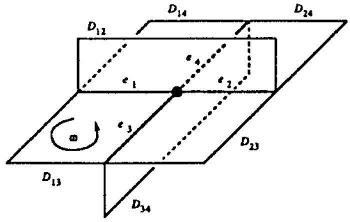

Figure 32: Lettering conventions used to describe the effect of C_I

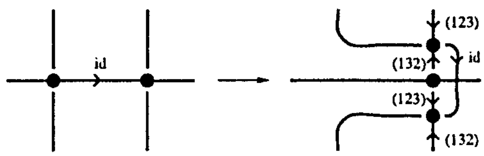

Figure 33: The move MP

situation of decorated standard polyhedra (*cf.* the proof of 2.6), but it is also not hard to prove it directly.

Proposition 2.8. *The move* MP *defined in Fig. 33 translates for s-graphs the Matveev-Piergallini move for an edge which has both the endpoints over and colour* id. *Moreover up to moves* R, C_I *and* C_{II} *every edge becomes like this.*

Proof of 2.8. The first assertion is proved by Figg. 34 and 35. The second assertion is easily deduced using Remark 2.7. 2.8

As a consequence of Theorems 1.1 and 1.2, Propositions 2.4, 2.6 and 2.8 we have:

Theorem 2.9. *For every compact* 3*-manifold* M *with non-empty boundary there exists an s-graph* Γ *with* $M = \mathcal{M}(\Gamma)$. *If* Γ_1 *and* Γ_2 *are s-graphs with at least two vertices such that* $\mathcal{M}(\Gamma_1)$ *and* $\mathcal{M}(\Gamma_2)$ *are defined (see Proposition 2.3), then* $\mathcal{M}(\Gamma_1)$ *and* $\mathcal{M}(\Gamma_2)$ *are homeomorphic if and only if* Γ_1 *and* Γ_2 *are obtained from each other by a sequence of moves* R, U, C_I, C_{II}, MP *and inverses of them.*

3 A graphic calculus for oriented 3-manifolds

In this section we specialize the construction of the previous one for oriented standard polyhedra, i.e. for standard spines of oriented manifolds. We first characterize the s-graphs which correspond to orientable manifolds.

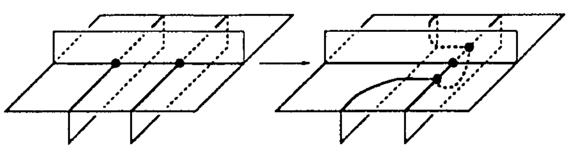

Figure 34: A representation of the Matveev-Piergallini move on polyhedra

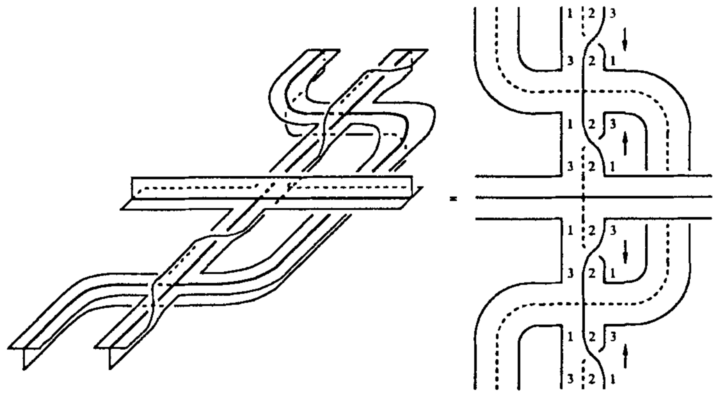

Figure 35: This shows the effect of the Matveev-Piergallini move on a neighbourhood of the singular part. The planar representation is immediately translated to the move MP

Proposition 3.1. *Let Γ be an s-graph. Then $\mathcal{P}(\Gamma)$ is orientable if and only if the product of the signs of colours along any simple loop is* $+1$.

Proof of 3.1. If one looks at the definition of the $\mathfrak{T}$-bundle over a circle which is a union of edges, one easily sees that each edge with negative sign induces a transposition in the fibres, while an edge with positive sign induces a rotation. Therefore, the bundle has two components if and only if the product of the signs is -1, and the conclusion follows from Theorem 1.5(B). $\boxed{3.1}$

We have stated this result as it could be useful when dealing with manifolds which are not a priori orientable. However we will show now that *oriented* manifolds admit a simpler representation. As in the previous section we start with an intrinsic definition of decoration and then we derive the graph representation.

Let (D,ω) be a decoration for a standard polyhedron P at a vertex v and denote as usual by π the union of the germs of discs adjacent to D. Let e be an edge of D, and direct e towards v. Let D' be the germ of disc on π on the right of e and D'' the germ of disc on π on the left of e (the notions of "right" and "left" involve ω and the direction of e). Now let P be oriented: we say the decoration is compatible with the orientation of P at v if the direction of e_1 and the cyclic order $D \to D' \to D'' \to D$ define the positive orientation of P along e_1. It is not hard to check that this is a natural choice.

Remark 3.2. Given an oriented standard polyhedron exactly 6 decorations at each vertex are compatible with the orientation.

Proposition 3.3. *Let Γ be an s-graph and let $\mathcal{D}(\Gamma)$ be the decorated standard polyhedron defined in Proposition 2.4. There exists an orientation compatible with the decoration at each vertex if an only if all the colours in Γ have sign* $+1$. *Such an orientation, if any, is unique.*

Proof of 3.3. We recall that Γ defines an embedding in $\mathbb{R}^3$ of a neighbourhood of each vertex, and the decoration is defined using these embeddings. Define the orientation

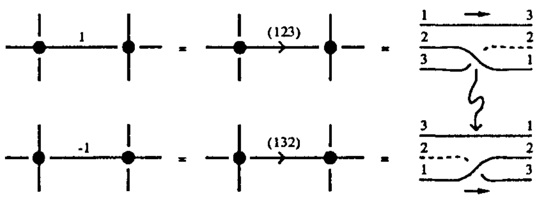

Figure 36: The reflection in the horizontal plane replaces a colour by its inverse

near every vertex so that compatibility is respected. Consider an edge: it is not hard to see that the orientations near its endpoints match along the edge if and only if the colour the edge bears has sign $+1$, and the conclusion easily follows. [3.3]

Recalling Remark 2.5 we have that for s-graphs as described in Proposition 3.3 the direction of the edges is actually dispensable, and the colours belong to the subgroup of $\mathfrak{S}_3$ generated by $(1\,2\,3)$, which is isomorphic to $\mathbb{Z}_3$. Hence if we define an o-graph as an s-graph where are edges are not directed and the colours are taken in $\mathbb{Z}_3$ rather than in $\mathfrak{S}_3$ (see the introduction for a self-contained definition) using Proposition 2.4 we have:

Proposition 3.4. *Decorated oriented standard polyhedra with decoration compatible with the orientation correspond bijectively to o-graphs up to isotopy and moves* R.

As for s-graphs, from now on we consider o-graphs up to isotopy and moves R. Given an o-graph we denote by $\mathcal{P}^{\Lambda}(\Gamma)$ the oriented standard polyhedron defined by Γ in which the decoration is forgotten.

Proposition 3.5. *Given o-graphs Γ_1 and Γ_2 there exists an orientation-preserving homeomorphism of $\mathcal{P}^{\Lambda}(\Gamma_1)$ onto $\mathcal{P}^{\Lambda}(\Gamma_2)$ if and only if Γ_1 and Γ_2 are obtained from each other by a finite sequence of applications of the move* C *defined in Fig. 3.*

Proof of 3.5. The move C is just the translation of C_I in the language of o-graphs. As in the proof of Proposition 2.6 we only need to check that C generates the changes of decoration compatible with the orientation. This can be left to the reader. [3.5]

Remark 3.6. There are two ways to apply the move C at every vertex, and exactly 6 different situations can be produced by repeated applications.

Proposition 3.7. *If Γ is an o-graph and $P^{\lambda} = \mathcal{P}^{\Lambda}(\Gamma)$ then $P^{-\lambda} = \mathcal{P}^{\Lambda}(\Gamma')$, where Γ' is obtained from Γ by reversing the under-over specification at marked vertices and replacing each colour by its inverse in $\mathbb{Z}_3$.*

Proof of 3.7. Γ defines an orientation-preserving embedding in $\mathbb{R}^3$ of a neighbourhood of the singular part of P. We prove that replacing Γ by Γ' corresponds to composing the embedding with the reflection of $\mathbb{R}^3$ in the horizontal plane. Of course the effect of the reflection on a vertex is to exchange under-arc and over-arc, while we show by an example in Fig. 36 that the effect on colours is inversion in $\mathbb{Z}_3$. [3.7]

Now, we have the move MP which is readily translated in the language of o-graphs, as shown in Fig. 4 of the introduction. We denote it by MP again.

Proposition 3.8. *The move* MP *translates for o-graphs the oriented Matveev-Piergallini move for an edge which has both the endpoints over and colour* 0. *Moreover up to moves* R *and* C *every edge becomes like this.*

Proof of 3.8. Recalling Proposition 2.8 the first assertion is obvious, and the second one is readily checked using Remark 3.6. [3.8]

We have eventually established the main result stated in the introduction. In fact Theorem 0.1 summarizes Theorems 1.1 and 1.2 and Propositions 3.4, 3.5, 3.7 and 3.8.

References

[1] R. BENEDETTI – C. PETRONIO, *On Roberts' proof of the Turaev-Walker theorem.* To appear in Journal of Knot Theory and its Ramifications

[2] B. G. CASLER, *An imbedding theorem for connected 3-manifolds with boundary.* Proc. Amer. Math. Soc. **16** (1965), 559-566

[3] R. A. FENN – C. P. ROURKE, *On Kirby's calculus on links.* Topology **18** (1979), 1-15

[4] R. KIRBY, *A calculus for framed links in S^3.* Invent. Math. **105** (1978), 35-46

[5] S. V. MATVEEV, *Special spines of piecewise linear manifolds.* Math. USSR-Sb. **21** (1973), 279-291

[6] S. V. MATVEEV, *Universal 3-deformations of special polyhedra.* Russ. Math. Surv. **42** (1987), 226-227

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

[8] C. PETRONIO, *Standard spines and their graphic presentation as a tool for studying 3-manifolds.* Tesi di Perfezionamento, Scuola Normale Superiore, Pisa, 1995

[9] R. PIERGALLINI, *Standard moves for standard polyhedra and spines.* Rendiconti Circ. Mat. Palermo **37**, suppl. 18 (1988), 391-414

[10] C. P. ROURKE – B. J. SANDERSON, An introduction to piecewise linear topology. Ergebn. der Math. Bd 69, Springer-Verlag, Berlin-Heidelberg-New York, 1982

[11] V. G. TURAEV – O. YA. VIRO, *State sum invariants of 3-manifolds and quantum 6j-symbols.* Topology **31** (1992), 865-902

Dipartimento di Matematica (Università di Pisa)
Via F. Buonarroti, 2 I-56127 PISA – Italy