

Evaluating TQFT invariants from G -crossed braided spherical fusion categories via Kirby diagrams with 3-handles

Manuel Bärenz

Abstract. A family of TQFTs parametrised by G -crossed braided spherical fusion categories has been defined recently as a state sum model and as a Hamiltonian lattice model. Concrete calculations of the resulting manifold invariants are scarce because of the combinatorial complexity of triangulations, if nothing else. Handle decompositions, and in particular Kirby diagrams are known to offer an economic and intuitive description of 4-manifolds. We show that if 3-handles are added to the picture, the state sum model can be conveniently redefined by translating Kirby diagrams into the graphical calculus of a G -crossed braided spherical fusion category.

This reformulation is very efficient for explicit calculations, and the manifold invariant is calculated for several examples. It is also shown that in most cases, the invariant is multiplicative under connected sum, which implies that it does not detect exotic smooth structures.

Contents

| | |
|---|----|
| 1. Introduction | 2 |
| 2. Introduction to G -crossed braided spherical fusion categories | 4 |
| 3. Kirby calculus with 3-handles | 7 |
| 4. Graphical calculus in G -crossed braided spherical fusion categories | 22 |
| 5. The invariant | 34 |
| 6. Calculations | 38 |
| 7. Recovering the state sum model | 44 |
| 8. Spherical fusion 2-categories and related work | 48 |
| A. Manifolds with corners | 50 |
| B. Details on Kirby calculus and diagrams | 52 |
| C. 1-handle slide lemmas in fusion categories | 54 |
| References | 56 |

2020 *Mathematics Subject Classification.* Primary 57K16; Secondary 57K41, 18M15, 18M20, 57R56, 18M30.

Keywords. Topological quantum field theory, fusion category, Kirby diagrams.

1. Introduction

In the study of 4-dimensional topological quantum field theories (TQFTs), few interesting examples are known, and even fewer are defined rigorously as axiomatic TQFTs. The fewest are economic to calculate.

Often, families of TQFTs like the Reshetikhin–Turaev model [25], the Turaev–Viro–Barrett–Westbury state sum [8, 31] or the Crane–Yetter model [10] are indexed by an algebraic datum such as a fusion category with extra structure. More recently [6, 32], G -crossed braided fusion categories have been studied in Hamiltonian approaches to topological phases, and a state sum model (and thus, an axiomatic TQFT) was defined [11]. This G -crossed model is thus defined explicitly, but it is hard to calculate concrete values for the partition function.

The Crane–Yetter model is famously invertible if the input datum is a modular category, but it had been suspected that it is non-invertible for nonmodular categories, with a rigorous proof only given more than two decades after its definition [5]. One reason seems to be that state sum models based on triangulations are easy to define, but hard to calculate. Once the partition function is defined in terms of handle decompositions, it is much easier to calculate concrete values, and indeed the second open question in [11, Section 7] asks for a description of the G -crossed model in terms of Kirby diagrams.

This is achieved here, adapting the techniques from [5]. As the central result, we define for every G -crossed braided spherical fusion category \mathcal{C} an invariant $I_{\mathcal{C}}$ of smooth, closed, oriented manifolds and show (Theorem 7.1):

Theorem. *Up to a factor involving the Euler characteristic, the invariant $I_{\mathcal{C}}$ is equal to the state sum $Z_{\mathcal{C}}$ in [11, (23)]. Explicitly, let M be a connected, smooth, oriented, closed 4-manifold and \mathcal{T} an arbitrary triangulation, then*

$$I_{\mathcal{C}}(M) = Z_{\mathcal{C}}(M; \mathcal{T}) \cdot d(\Omega_{\mathcal{C}})^{1-\chi(M)} \cdot |G|.$$

The invariant $I_{\mathcal{C}}$ is in fact much easier to calculate than the state sum, and several examples are given. It is also shown that, in most cases, $I_{\mathcal{C}}$ does not differentiate between smooth structures on the same topological manifold.

Triangulations and handle decompositions. As a general theme in many flavours of topology, spaces are built up from simple elementary building blocks. A large class of topological spaces can be constructed as simplicial complexes or as CW-complexes. In the former approach, the elementary building blocks are simpler and the possibilities of gluing them are fewer, whereas in the latter, spaces can often be described much more succinctly. In both cases, each building block has an inherent dimension, which is a natural number. Spaces with building blocks of dimension at most n are usually said to be n -dimensional.

For smooth manifolds, the situation is similar. They admit triangulations, which yield simplicial complexes, but also handle decompositions, which are analogous to CW-complexes. Triangulations decompose an n -manifold into k -simplices, and handle decompositions consist of k -handles, respectively for $0 \leq k \leq n$.

We favour handle decompositions over triangulations here because they are much more succinct.¹ One can specify the 4-dimensional sphere S^4 with as little as a 0-handle and a 4-handle, while the standard triangulation coming from the boundary of the 5-simplex contains 20 triangles. It is no surprise that state sum models are often defined for triangulations, but no values are calculated even for simple manifolds (as in [10, 11]). Evaluating TQFTs on bordisms via handle decompositions is both conceptually more direct [19] and computationally more efficient. For example, the Crane–Yetter model was famously shown to be invertible for modular categories using handle decompositions [26, Theorem 4.5], and most concrete values for pre-modular categories have only been calculated more than 20 years after its state sum definition, again using handle decompositions [5].

The price to pay is the additional complexity when attaching handles to each other. While the gluing data of simplices is completely combinatorial, handle attachment data is quite topological in nature, and is best described in diagrams. In contrast to simplices, handles of all dimensions $0 \leq k \leq n$ must be thickened to n dimensions before they are glued along their $(n - 1)$ -dimensional boundaries. Consequently, all of the attachment data can be described diagrammatically in $(n - 1)$ -dimensions. We thus describe 4-manifolds with 3-dimensional diagrams.

To evaluate the TQFT on them, these diagrams are decorated with data from a (non-strict) 3-category. In particular, a G -crossed braided category is a special case of a monoidal 2-category [11, Section 6], and thus a special case of a 3-category. And indeed, there is a beautiful way to label diagrams of 4-dimensional handle decompositions with G -crossed braided spherical fusion categories ($G \times$ -BSFCs), by depicting 3-handles explicitly in the calculus.

Outline. We begin by briefly introducing G -crossed braided spherical fusion categories in Section 2. Additionally to defining the invariant, this article gives an introduction to Kirby diagrams with 3-handles, which does not seem to exist in the literature. Section 3 contains this material, and it should be accessible to TQFT researchers with a basic background in Kirby calculus of 4-manifolds. The aim of Section 4 is to show

¹In higher dimensions, there is another reason: triangulations do not completely specify a smooth structure, but only a piecewise linear (PL) structure, whereas handle decompositions exist both in the smooth and in the PL category. In four dimensions, the PL and smooth categories are still equivalent, which allows us to formulate smooth TQFTs as state sum models over triangulations.

that the language of Kirby diagrams with 3-handles is very natural for graphical calculus in G -crossed braided spherical fusion categories. This is the central ingredient to Section 5, where the invariant $I_{\mathcal{C}}$ is defined. Explicit example calculations and general properties of the invariant are given in Section 6. The connection to the state sum model from [11] and the corresponding TQFT is made in Section 7. Further ideas for generalisations, in the direction of defining spherical fusion 2-categories, are discussed in Section 8.

2. Introduction to G -crossed braided spherical fusion categories

2.1. Spherical fusion categories

Spherical fusion categories over \mathbb{C} are well known, and basic knowledge is assumed. We will mainly use notation, conventions and graphical calculus from [5, Section 2.1]. For convenience, an overview over the relevant notation is given in Table 1. We will augment the graphical calculus of fusion categories with “round coupons” for morphisms (this is justified because a pivotal structure is assumed to exist) as introduced in [3, Section 5.3] and explained for example in [4, Section 1]. Essentially, morphisms do not differentiate between source and target, as these can be interchanged coherently by means of duals.

| | |
|---|--|
| Spherical fusion category | $(\mathcal{C}, \otimes, I)$ |
| Objects | A, B, C, \dots |
| Morphisms | $\iota \in \langle A \otimes B \otimes \dots \rangle := \mathcal{C}(I, A \otimes B \otimes \dots)$ |
| Simple objects | X, Y, Z |
| Set of representants of isomorphism classes of simple objects | $\mathcal{O}(\mathcal{C})$ |
| Dual object | A^* |
| Nondegenerate morphism pairing | $(-, -): \langle A_1 \otimes \dots \otimes A_n \rangle \otimes \langle A_n^* \otimes \dots \otimes A_1^* \rangle \rightarrow \mathbb{C}$ |
| Trace | $\text{tr}: \mathcal{C}(A, A) \rightarrow \mathbb{C}$ |
| Categorical dimension | $d(X) := \text{tr}(1_X)$ |
| Fusion algebra (complexified Grothendieck ring) | $\mathbb{C}[\mathcal{C}]$ |
| “Kirby colour” | $\Omega_{\mathcal{C}} := \bigoplus_{X \in \mathcal{O}(\mathcal{C})} d(X)X \in \mathbb{C}[\mathcal{C}]$ |

Table 1. Notation and conventions in spherical fusion categories.

We briefly revisit the cyclically symmetric definition of morphism spaces:

$$\begin{aligned}
 \iota \in \langle A_1 \otimes A_2 \otimes \cdots \otimes A_n \rangle &:= \mathcal{C}(\mathcal{I}, A_1 \otimes A_2 \otimes \cdots) \\
 &\stackrel{\text{pivotal structure}}{\cong} \mathcal{C}(\mathcal{I}, A_{j+1} \otimes \cdots \otimes A_n \otimes A_1 \otimes \cdots \otimes A_j) \\
 &\stackrel{\text{dualisation}}{\cong} \mathcal{C}(A_j^* \otimes A_{j-1}^* \otimes \cdots \otimes A_1^* \otimes A_n^* \otimes \cdots \otimes A_k^*, A_j \otimes \cdots \otimes A_{k-1}).
 \end{aligned}$$

The morphism pairing $(-, -): \langle A_1 \otimes \cdots \otimes A_n \rangle \otimes \langle A_n^* \otimes \cdots \otimes A_1^* \rangle \rightarrow \mathbb{C}$ is defined by dualisation, composition, and the isomorphism $\mathcal{C}(\mathcal{I}, \mathcal{I}) \cong \mathbb{C}$. It is *nondegenerate*, which implies the existence of dual bases

$$\{\alpha_i\} \subset \langle A_1 \otimes \cdots \otimes A_n \rangle \quad \text{and} \quad \{\tilde{\alpha}_i\} \subset \langle A_n^* \otimes \cdots \otimes A_1^* \rangle$$

satisfying $(\tilde{\alpha}_i, \alpha_j) = \delta_{i,j}$. The choice of $\{\alpha_i\}$ is arbitrary and determines $\{\tilde{\alpha}_i\}$ completely.

Definition 2.1. Together, the dual bases define a unique entangled vector

$$\sum_i \alpha_i \otimes \tilde{\alpha}_i \in \langle A_1 \otimes \cdots \otimes A_n \rangle \otimes \langle A_n^* \otimes \cdots \otimes A_1^* \rangle.$$

Especially in the graphical calculus (Figure 1), this vector will often be abbreviated as $\alpha \otimes \tilde{\alpha}$, suppressing the indices.

$$\begin{array}{c}
 \begin{array}{ccc}
 \begin{array}{c} A_2 \\ \diagdown \\ \bigcirc \\ \diagup \\ A_1 \end{array} & & \begin{array}{c} A_1^* \\ \diagdown \\ \bigcirc \\ \diagup \\ A_2^* \end{array} \\
 \begin{array}{c} A_3 \\ \diagdown \\ \bigcirc \\ \diagup \\ A_4 \end{array} & \alpha & \begin{array}{c} A_6 \\ \diagdown \\ \bigcirc \\ \diagup \\ A_5 \end{array} \\
 \end{array} \quad
 \begin{array}{ccc}
 \begin{array}{c} A_6^* \\ \diagdown \\ \bigcirc \\ \diagup \\ A_3^* \end{array} & \tilde{\alpha} & \begin{array}{c} A_1^* \\ \diagdown \\ \bigcirc \\ \diagup \\ A_2^* \end{array} \\
 \begin{array}{c} A_5^* \\ \diagdown \\ \bigcirc \\ \diagup \\ A_4^* \end{array} & & \begin{array}{c} A_1^* \\ \diagdown \\ \bigcirc \\ \diagup \\ A_2^* \end{array} \\
 \end{array}
 \end{array}
 = \sum_i
 \begin{array}{ccc}
 \begin{array}{c} A_2 \\ \diagdown \\ \bigcirc \\ \diagup \\ A_1 \end{array} & \alpha_i & \begin{array}{c} A_6 \\ \diagdown \\ \bigcirc \\ \diagup \\ A_5 \end{array} \\
 \begin{array}{c} A_3 \\ \diagdown \\ \bigcirc \\ \diagup \\ A_4 \end{array} & & \begin{array}{c} A_1^* \\ \diagdown \\ \bigcirc \\ \diagup \\ A_2^* \end{array} \\
 \end{array}
 \quad
 \begin{array}{ccc}
 \begin{array}{c} A_6^* \\ \diagdown \\ \bigcirc \\ \diagup \\ A_3^* \end{array} & \tilde{\alpha}_i & \begin{array}{c} A_1^* \\ \diagdown \\ \bigcirc \\ \diagup \\ A_2^* \end{array} \\
 \begin{array}{c} A_5^* \\ \diagdown \\ \bigcirc \\ \diagup \\ A_4^* \end{array} & & \begin{array}{c} A_1^* \\ \diagdown \\ \bigcirc \\ \diagup \\ A_2^* \end{array} \\
 \end{array}
 \end{array}$$

Figure 1. The dual bases convention with round coupons. Compare, e.g., [4, (1.8)].

2.2. G -crossed braided spherical fusion categories

G -crossed braided (spherical) fusion categories (short: $G \times$ -BSFCs) were introduced in [29] as “crossed group categories” and have since received a diversity of names such as G -crossed braided SFCs, G -equivariant braided SFCs or braided G -crossed SFCs. We will adopt the nomenclature $G \times$ -BSFC, but warn that $G \times$ -BSFCs are usually *not* braided. Rather, they carry a new structure, the *crossed braiding*.

Just as fusion categories categorify and vastly generalise finite groups, $G \times$ -BSFCs categorify finite crossed modules. (This viewpoint is implemented explicitly in [11,

Section 4.2].) A crossed module consists of two groups H and G , a homomorphism $\text{deg}: H \rightarrow G$ and a group action of G on H . The group H is “crossed commutative,” that is, commutative up to an action by G . This axiom is usually called the “Peiffer rule.” $G \times$ -BSFCs now categorify finite crossed modules in the following way. The group H is generalised to a spherical fusion category \mathcal{C} , while G stays a group. The homomorphism deg gives way to a G -grading of \mathcal{C} . Just as abelian groups are usually categorified to braided fusion categories and the commutativity axiom is replaced by the braiding, the Peiffer rule is now replaced by a crossed braiding.

Writing out the axioms explicitly and following the notation from [11, Section 2.1], we have the following definitions.

Definitions 2.2. Let G be a finite group and \mathcal{C} a spherical fusion category. Denote by \underline{G} the discrete monoidal category whose objects are elements of G and the monoidal structure matches the group structure.

- A G -grading on \mathcal{C} is a decomposition $\mathcal{C} \cong \bigoplus_{g \in G} \mathcal{C}_g$ into semisimple linear categories such that $\mathcal{C}_{g_1} \otimes \mathcal{C}_{g_2} \subset \mathcal{C}_{g_1 g_2}$. This defines a function $\text{deg}: \mathcal{O}(\mathcal{C}) \rightarrow G$.
- A G -action on \mathcal{C} is a monoidal functor $(F, \eta, \varepsilon): \underline{G} \rightarrow \text{Aut}_{\otimes, \text{piv}}(\mathcal{C})$. Explicitly, there is a monoidal, pivotal automorphism of \mathcal{C} for each group element g . We will usually abbreviate ${}^g X := F(g)(X)$, and thus the coherence isomorphisms are $\eta(g_1, g_2)_X: {}^{g_1}({}^{g_2} X) \xrightarrow{\cong} {}^{g_1 g_2} X$ and $\varepsilon_X: {}^e X \xrightarrow{\cong} X$.
- A G -crossed braided spherical fusion category consists of the following structure and axioms:
 - a spherical fusion category \mathcal{C} ,
 - a G -grading on \mathcal{C} ,
 - a G -action on \mathcal{C} such that ${}^g(\mathcal{C}_{g'}) \subset \mathcal{C}_{g g' g^{-1}}$,
 - for every $g \in G, X \in \mathcal{C}_g, Y \in \mathcal{C}$, a natural isomorphism $c_{X, Y}: X \otimes Y \rightarrow {}^g Y \otimes X$, which is called the *crossed braiding*,
 - subject to certain axioms which are detailed, e.g., in [11, Definition 2.2] or [13, Definition 4.41].

Remark 2.3. A $G \times$ -BSFC is usually not braided. c is only a braiding if the G -action is trivial.

Remark 2.4. The G -grading is not required to be faithful, i.e., \mathcal{C}_g may be 0 for some $g \neq e$. In contrast to mere graded fusion categories, such a $G \times$ -BSFC is not always equivalent to a faithfully graded one for a subgroup of G , since the G -action can contain information about the whole group.

Examples 2.5. (1) Any ribbon fusion (“premodular”) category is a spherical braided fusion category, and thus a $G \times$ -BSFC for G the universal grading group and the trivial action.

(2) [13, Section 4.4] explains how a braided inclusion $\text{Rep}(G) \subset \mathcal{D}$ of the representations of a finite group into a braided fusion category yields a $G \times$ -BSFC.

3. Kirby calculus with 3-handles

3.1. Handle decompositions

In this section, handle decompositions of 4-manifolds and their 3-dimensional diagrams are described. (All manifolds will be assumed to be smooth, oriented and compact.) While handle decompositions and Kirby diagrams are described extensively in the literature [1, 15, 21], 3-handles are rarely described pictorially, the only case known to the author being [16]. This will turn out to be an important conceptual clarification and computational simplification in the graphical calculus of $G \times$ -BSFCs, and is thus introduced here at length, interspersed with a recapitulation of well-known material.

Definitions 3.1. An n -dimensional k -handle is the manifold with corners² $h_k := D^k \times D^{n-k}$.

- From its corner structure, the boundary of a k -handle $\partial h_k = S^{k-1} \times D^{n-k} \cup D^k \times S^{n-k-1}$ is split into the *attaching region* $\partial_a h_k = S^{k-1} \times D^{n-k}$ and the *remaining region* $\partial_r h_k = D^k \times S^{n-k-1}$.
- $S^{k-1} \times \{0\} \subset \partial_a h_k$ is called the *attaching sphere*.
- $\{0\} \times S^{n-k-1} \subset \partial_r h_k$ is called the *remaining sphere*, or *belt sphere*.

We wish to decompose n -manifolds M^n into a filtration $\emptyset = M_{-1} \subset M_0 \subset M_1 \subset \dots \subset M_n \cong M$, where each M_k is produced from M_{k-1} by attaching k -handles. This decomposition is then called a *handlebody*, and it is often used interchangeably with M , or regarded as extra structure on M .

Definitions 3.2. • A k -handle attachment map ϕ on an n -manifold M is an embedding of the attaching region $\partial_a h_k$ of an n -dimensional k -handle into ∂M .

- The result of a handle attachment $\phi: \partial_a h_k \hookrightarrow \partial M$ is the manifold $M \cup_\phi h_k$, where the handle is glued along the embedded attaching region.
- A (-1) -handlebody is the empty set.

²Details and references about manifolds with corners and their boundaries can be found in Appendix A.

- A k -handlebody is a $(k - 1)$ -handlebody and successive k -handle attachments on it.
- A handle decomposition of a manifold M^n is a diffeomorphism to an n -handlebody.

Remark 3.3. When attaching a k -handle to a manifold M along ϕ , the attaching boundary is removed and the remaining boundary added. The boundary of the result of a handle attachment is thus

$$\partial(M \cup_{\phi} h_k) = \partial M \setminus \phi(\partial_a h_k) \cup \partial_r h_k.$$

This operation is known as *performing surgery* on ∂M along ϕ .

Remark 3.4. By definition, the attaching regions of the attached handles are in the interior of the resulting manifold, therefore further handles are always attached to the remaining regions of previous handles, explaining the naming choice “remaining.”

Theorem 3.1 (Well known, e.g., [15, Section 4.2]). *Smooth manifolds have handle decompositions. In particular, smooth compact manifolds have decompositions with finitely many handles.*

The proof of this theorem is usually via Morse theory. Given a self-indexing Morse function, every critical point of index k corresponds to a k -handle.

Since arbitrarily many different handle decompositions may describe the same manifold up to diffeomorphism, it is important to relate them. Luckily, there is a simple complete set of moves to translate from any two diffeomorphic handle decompositions.

Definition 3.5. A k -handle h_k and a $(k + 1)$ -handle h_{k+1} are *cancellable* if the attaching sphere of h_{k+1} intersects the remaining sphere (the belt sphere) transversely in one point.

The terminology stems from the fact that $M \cup h_k \cup h_{k+1} \cong M$ if h_k and h_{k+1} are cancellable. This diffeomorphism is called *cancelling* the handle pair.

Theorem 3.2 (E.g., [15, Theorem 4.2.12]). *Two handle decompositions of the same manifold are related by a finite sequence of handle attachment map isotopies, attachment sequence reorderings, and handle pair cancellations.*

Remark 3.6. A k -handle can be isotoped over another k -handle in a canonical way, this is called a *handle slide* [15, Definition 4.2.10]. Since the attaching order for handles of the same level can be changed arbitrarily, any k -handle can be slid over any other k -handle.

Theorem 3.3 (E.g., [15, Proposition 4.2.13]). *A connected, closed, smooth n -manifold has a handle decomposition with exactly one 0-handle and exactly one n -handle.*

The idea of the proof is to cancel all 0-1-pairs until a single 0-handle remains, and dually for the n -handle.

In general, it is not clear how to visualise intricate handle attachment maps. In four dimensions though, all attachments happen inside three-dimensional spaces, which allows us to depict them diagrammatically.

3.2. Kirby diagrams

For the remainder of the article, we will assume handle decompositions to be in dimension $n = 4$. The relevant special cases for 4-dimensional handles and their boundaries are listed in Table 2.

| k | Handle h_k | Attaching region $\partial_a h_k$ | Remaining region $\partial_r h_k$ |
|-----|------------------|---|---|
| 0 | $D^0 \times D^4$ | \emptyset | $S^3 \cong \mathbb{R}^3 \cup \{\infty\}$ |
| 1 | $D^1 \times D^3$ | $S^0 \times D^3 \cong \{-1, 1\} \times D^3$ | $D^1 \times S^2 \cong [-1, 1] \times S^2$ |
| 2 | $D^2 \times D^2$ | $S^1 \times D^2$ | $D^2 \times S^1$ |
| 3 | $D^3 \times D^1$ | $S^2 \times D^1$ | $D^3 \times S^0$ |

Table 2. 4-dimensional k -handles and their boundary, for $k \leq 3$.

3.2.1. Remaining regions as canvases. It is possible to visualise handle decompositions of 4-manifolds as Kirby diagrams by thinking of the remaining regions of the handles as *drawing canvases*. In fact, the remaining region of a 0-handle can be visualised as \mathbb{R}^3 with an additional point at infinity (Figure 2). Since $\partial_a h_k \cup \partial_r h_k \cong S^3$, the remaining regions of higher handles can be visualised in the same way, although with the attaching region removed from the canvas.

Initially, it is always possible to attach a 0-handle to the empty manifold, which instantiates a new, empty drawing canvas. Higher handles for $1 \leq k \leq 3$ are attached to the existing handlebody by drawing their attaching region (Figure 3) on the existing canvases, which corresponds to embedding it onto the boundary of the handlebody. Here, it will be enough to simply draw the attaching *sphere*, as will be justified shortly.

As in Theorem 3.3, a 1-handle which is attached to two different 0-handles can cancel one of them (Figure 4), merging the two canvases. If further handles are attached to the cancelling 0-handle, their attachment maps can be isotoped across the (remaining boundary of the) 1-handle into the (remaining boundary of the) other 0-handle. We will therefore usually assume handle decompositions to have exactly one 0-handle, and omit the coordinate axes. (By the same theorem, we will assume exactly one 4-handle to be present, and never explicitly specify it.) Examples of higher cancellations and slides will be reviewed in Section 3.3.




Figure 2. The remaining region of a 0-handle as a drawing canvas.

| Handle index k | Attaching region $\partial_a h_k$ |
|------------------|-----------------------------------|
| 0 | \emptyset |
| 1 | |
| 2 | |
| 3 | |

Figure 3. Attaching regions of 1-, 2- and 3-handles.

A handle can be attached onto several other handles of smaller index, in which case its attaching sphere is spread over several remaining regions, as for example in Figure 5.

3.2.2. Attaching spheres and framings. Up to isotopy, 1-handle and 3-handle attachments are determined by the embedding of the attaching sphere: an embedding of $\partial_a h_1 = S^0 \times D^3$ is, up to isotopy, specified by the embedding of the two points of S^0 .

Similarly, a 3-handle attachment is essentially specified by the embedding of an S^2 [15, Example 4.1.4 c]. Since usually further handles are attached to 1-handles, but




Figure 4. The handles $h_{0,B}$ and $h_{1,\beta}$ cancel each other. The 0-1-cancellation move merges two drawing canvases.




Figure 5. A handle decomposition of $[-1, 1] \times \mathbb{R}P^3$ where each k -handle is attached to every j -handle for $j < k$. In particular, the single 3-handle is attached to h_1 and h_2 : its attaching sphere touches the attaching spheres of h_1 and h_2 in ∂h_0 (and h_2 's attaching sphere further in $\partial_r h_1$), and enters their respective remaining regions, partitioning it into six pieces.

not to 3-handles, we will draw the complete attaching regions of 1-handles, but only the attaching spheres of 3-handles, as in Figure 3.

Furthermore, it is sometimes helpful to mark the orientation of S^0 in $\partial_a h_1 = S^0 \times D^3$ by denoting one point as $+$ and the other as $-$, and consequently denoting $\partial_a h_1 \cong \partial_a^+ h_1 \sqcup \partial_a^- h_1 = D_+^3 \sqcup D_-^3$.³ Similarly, the attaching S^2 of a 3-handle may carry an orientation. The result of the attachment does not depend on the choice of orientation.

A 2-handle is attached along $\partial_a h_2 = S^1 \times D^2$, and an embedding $S^1 \hookrightarrow \mathbb{R}^3$ can be knotted, or even linked to other 2-handles. The attachment carries more information than the mere embedding of its attaching S^1 , though. It may be *framed*, that is, the D^2 component can be twisted by any integer multiple of a full turn. It is possible to denote framings by labelling 2-handle attachments with integers, but for our purposes it is more pragmatic to stipulate the *blackboard conventions*: our diagrams are not truly drawn in \mathbb{R}^3 , but in a projection onto \mathbb{R}^2 , and this projection specifies a canonical framing on every curve. For further details, we refer to Appendix B.2.

3.2.3. Kirby conventions. Drawing the attaching spheres inside all remaining regions is uneconomical, but fortunately not necessary. Assume we have already attached a k -handle h_k to a manifold ($k \in \{1, 2\}$), and then attach a j -handle h_j for $j > k$ gradually, starting to draw the attaching sphere in $\partial_r h_0$ and aiming to continue into $\partial_r h_k$. Then the attaching sphere of h_j will eventually intersect the boundary of the attaching region of h_k , $\partial \partial_a h_k \cong S^{k-1} \times S^{3-k}$. This intersection already canonically determines an attachment inside $\partial_a h_k$, as we will see shortly. By always assuming this convention, the part of the diagram inside $\partial_r h_0$ completely determines the handle attachment.

Definition 3.7. A handle decomposition is *regular* if

- (1) the attachment of a k -handle is outside the remaining regions of other k -handles;
- (2) no attachment intersects with $\infty \in \partial_r h_0 \cong \mathbb{R}^3 \cup \{\infty\}$.

Remarks 3.8. (1) This condition is equivalent to the requirement k -handles are attached only onto the remaining regions of j -handles for $j < k$ (strictly).

(2) This condition ensures that the diagram can be drawn in a bounded region of \mathbb{R}^3 .

Definition 3.9. A handle decomposition satisfies the *single-picture conventions* if

- 2-1 inside the remaining region of any 1-handle $\partial_r h_1 = D^1 \times S^2 = [-1, 1] \times S^2$, images of attaching maps of 2-handles are of the form $[-1, 1] \times \{p_1, \dots, p_N\}$ with $p_i \in S^2$;

³This is not to be confused with the independent notion of relative handle decompositions.

3-1 inside the remaining region of any 1-handle $\partial_r h_1 = D^1 \times S^2 = [-1, 1] \times S^2$, images of attaching maps of 3-handles are of the form $[-1, 1] \times A$ with $A \subset S^2$ a compact 1-dimensional submanifold;

3-2 inside the remaining region of any 2-handle $\partial_r h_2 = D^2 \times S^1$, images of attaching maps of 3-handles are of the form $D^2 \times \{p_1, p_2, \dots, p_N\}$, where $p_i \in S^1$.

Remark 3.10. These two conditions have intuitive geometrical interpretations, on which we will expand.

2-1 A 2-handle entering one attaching ball D^3_+ of a 1-handle at a set of points must leave the corresponding ball D^3_- on the mirror positions. This is a standard assumption which is usually made in text books.

3-1 Analogously, a 3-handle entering one attaching ball D^3_+ of a 1-handle at a submanifold A must leave the corresponding ball D^3_- on the reflection of A .

3-2 A 3-handle enters a 2-handle along $S^1 \times \{p_1, p_2, \dots, p_N\} \subset S^1 \times S^1$, where each $S^1 \times \{p_i\}$ follows the framing of the 2-handle attachment. To the author's knowledge, this convention is unmentioned in the literature, although it is straightforward.

Definition 3.11. A regular handle attachment satisfying the single-picture conventions is called a *Kirby diagram*.

It is natural to ask whether essentially every handle decomposition can be drawn as a Kirby diagram. This is ensured by the following lemma.

Lemma 3.12. *Given any handle decomposition, a succession of isotopies of the individual handle attachments can be applied to arrive at a Kirby diagram.*

Proof. When disregarding 3-handles, this is a standard fact. See Figure 6 for an illustration. The full proof, covering the 3-handle case, is given in Appendix B. ■

3.2.4. Kirby diagrams and 3-handles. Usually, 3-handles are not depicted in Kirby diagrams, due to the following theorem.

Theorem 3.4 (Well known, e.g., [15, Section 4.4]). *Let M and N be smooth, closed, oriented 4-manifolds, with handle decompositions containing precisely one 0-handle and one 4-handle. Assume that their 2-handlebodies agree. Then $M \cong N$.*

In essence, the 3-handle attachments contain no information in a handle decomposition of a closed manifold. 4-dimensional differential topology usually concerns itself with closed manifolds, and 3-handles are typically excluded from most discussions. Standard references for Kirby diagrams such as [15] and [1] emphasise 1-handle and 2-handle attachments on a single 0-handle.

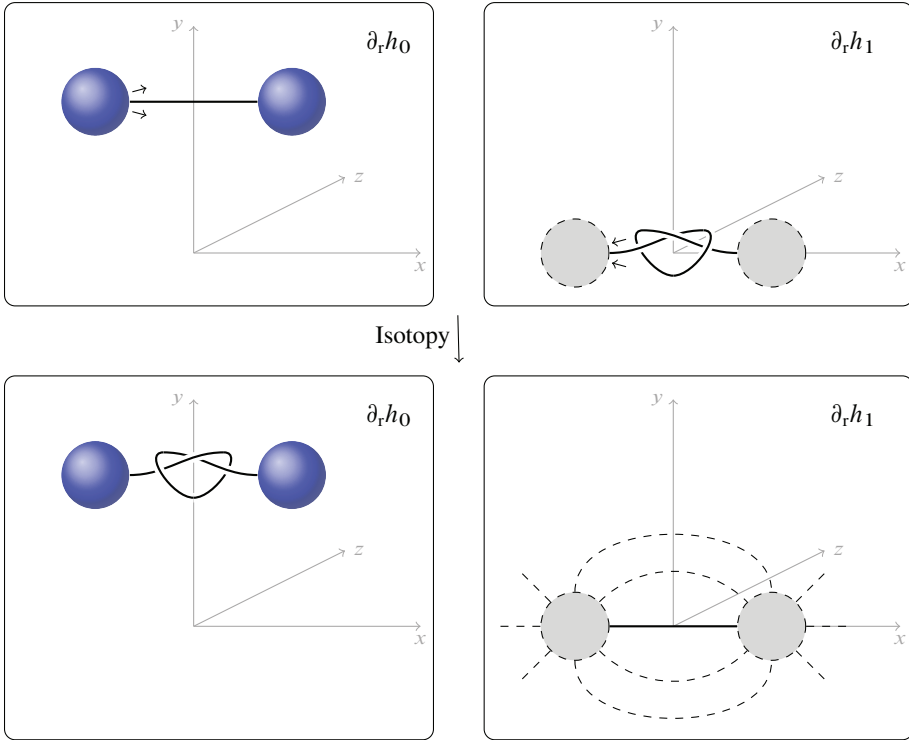


Figure 6. Pushing a 2-handle attachment outside the 1-handle remaining region. After the isotopy, the 2-handle attaching sphere runs along a canonical interval inside of $\partial_r h_2$. (Other canonical intervals are drawn dashed.) It is then easy to see that the 2-handle cancels the 1-handle, since it intersects the belt sphere (the $y - z$ -plane compactified at ∞) transversely in one point.

While it is true that the 2-handles contain the main complexity of a handle decomposition, 3-handles will still turn out to contain useful information for the computation of the invariant defined in Section 5.

As a further reason whose ramifications are beyond the scope of this article, Theorem 3.4 does not hold for manifolds with nonempty boundary. When our aim is to describe the whole TQFT directly with handle decompositions (without taking the detour over the state sum model), we will eventually have to regard bordisms, and there are indeed non-diffeomorphic manifolds with boundary that have handle decompositions with the same 2-handlebody⁴ Thus, 3-handles may indeed contain relevant information in this situation. This viewpoint is discussed briefly in Section 7.2.

⁴The author is grateful to Marco Golla and Andrew Lobb to point this fact out here: <https://mathoverflow.net/questions/288246/are-there-kirby-diagrams-with-3-handles>

3.3. Handle moves

Handle decompositions are by far not unique. Fortunately, Theorem 3.2 ensures that every two finite handle decompositions of the same closed manifold are related by a finite sequence of handle moves. There are three kinds of moves to consider: (1) handle cancellations, (2) (ambient) isotopies of handle attachments; (3) reorderings of attachments of the same index. Moves of the kind (1) are described below in Section 3.3.1. Moves of the kind (2) and of the kind (3) are described in Section 3.3.2.

From here on, we will not draw the coordinate grids of the remaining regions anymore, and assume that all diagrams take place in $\partial_r h_0$ (unless specified otherwise).

3.3.1. Cancellations. In our Kirby diagrams, cancellable pairs (Definition 3.5) of a k - and a $(k + 1)$ -handle are relevant for $k \in \{1, 2\}$. (We have described 0-1-cancellation already in Figure 4 in Section 3.2.1, and assume that there is a single 4-handle.)

A 2-handle cancels a 1-handle if the attaching circle attaches in exactly one point for each 3-ball of the 1-handle. This can be seen in Figure 6, where it was argued that any such attachment can be isotoped to one that intersects the belt sphere in one point. If further handles are attached to the handle pair, they can be slid off first before cancelling [15, Figure 5.13].

A 3-handle cancels a 2-handle if the 2-handle is attached along an unframed unknot, and (the visible part of the) 3-handle attachment forms a disc which is bounded by said unknot. (The second half of the attaching S^2 is inside the remaining region of the 2-handle, where it intersects the belt sphere in one point.) The situation can be seen in Figure 7. As for the 1-2-cancellation, it is sometimes necessary to first slide off other handles before cancelling.




Figure 7. The visible part of the 2-3-handle cancellation. The 3-handle attaching sphere is split into two discs, one of which vanishes inside the 2-handle remaining region.

3.3.2. Slides. Applying an isotopy to a constituent of a Kirby diagram (tautologically) applies an isotopy to the corresponding handle attachment, but there may be isotopies of handle attachments that do not conform to the Kirby conventions in the intermediate stages of the isotopy. Indeed, isotopies of handle attachment between two non-isotopic Kirby diagrams exist. They are commonly called “handle slides,” and they can be decomposed into elementary j - k -slides where $j \geq k$. A j - k -slide is an isotopy of a j -handle over the remaining region of a k -handle that changes the Kirby diagram.

If an l -handle is attached to a j -handle, the j - k -slide can still be performed simultaneously with an l - k -slide.

In order to allow for arbitrary k - k -slides, the attachment order of k -handle attachments has to be changed sometimes. Subtly, Kirby diagrams do not specify the order since the attachments always commute due to the first regularity condition (Definition 3.7), and thus resulting manifolds for different orders are diffeomorphic. When isotoping a k -handle h_k over another, the latter has to be attached first, thus a reordering may be necessary.

j - k -slides for $j \leq 2$ are treated at length in the literature, see, e.g., [15, Section 5.1]. For completeness, they are briefly described in the following, enumerated as j - k .

j -1 To slide any handle h_j over a 1-handle h_1 , choose a path from the attached $\partial_a h_j$ to $\partial_a h_1$ that does not intersect with any other attachments (in particular no 3-handle attachments) and move h_j 's attaching sphere along into $\partial_r h_1$, entering at one of the S^2 's bounding an attaching ball of h_1 .

Since the two attaching balls D_+^3 and D_-^3 can be situated in two regions of the drawing canvas that are separated by a 3-handle attachment, it is possible to move attachments between separated regions by means of this move.

1-1 Here, regularity requires the attachment to move all the way through the remaining boundary, leaving the opposite attaching ball again.

2-1 To establish the Kirby conventions, the 2-handle attachment needs not move completely through $\partial_r h_1$, instead its ends can protrude from the attaching balls of h_1 at mirror positions. The situation is displayed in Figure 5.

2-2 The pair of pants surface is defined as a two-dimensional disc D_c^2 with two open discs $D_a^{2,\circ}$ and $D_b^{2,\circ}$ removed, its boundary is thus $S_a^1 \sqcup S_b^1 \sqcup S_c^1$. Given an embedding of the pair of pants in a Kirby diagram K , its boundary embedding defines three 2-handle attachments $h_{2,a}$, $h_{2,b}$ and $h_{2,c}$. The 2-2-slide transforms K together with the attachments $h_{2,a}$ and $h_{2,b}$ into K with $h_{2,c}$ and $h_{2,b}$, and one says that $h_{2,a}$ has been slid over $h_{2,b}$.

For any two 2-handles, it is always possible to slide one over the other. (If they are in different regions separated by 3-handles, it is still possible to perform the 2-2-slide after a series of 2-1-slides, assuming the manifold is closed.) The resulting handle $h_{2,c}$ may depend on the choice of the pair of pants, but the resulting manifold does not (up to diffeomorphism).

Figures 8 and 9 display these slides, but also include analogous 3- k -slides, which are described in the following.




Figure 8. 1-1-handle slide with simultaneous 2-1-slide and 3-1-slide. The red 3-handle attachment separates the two components $\partial_a h_1 \cong D^{3+} \sqcup D^{3-}$ of the 1-handle attachments, but D_α^{3+} can still “tunnel” through the remaining region of $h_{1\beta}$.




Figure 9. 2-2-handle slide of h_{2a} over h_{2b} , simultaneously with a 3-2-slide of h_{3A} over h_{2b} . The grey area can contain arbitrary handles, and may share 3-handles with h_{2b} . Furthermore, h_{2b} may be arbitrarily knotted, and linked to other 2-handles. h_{2a} then follows those knots and links.

3-1 As in any j -1-slide, a 3-handle attachment can be isotoped through the remaining region of a 1-handle. And as in the 2-1-slide, the attachment can protrude from the attaching balls of the 1-handle at mirror positions.

For an explicit model of this situation, realise that the attaching sphere of the 3-handle can be split in three parts, consisting of two caps and a cylinder:

$$S^2 \cong D_+^2 \cup_{S_+^1} (S^1 \times [+1, -1]) \cup_{S_-^1} D_{2,-}^2.$$

Choose an embedding $S^1 \subset \partial D_+^3$ on the boundary of one ball of the 1-handle attachment. This induces a mirror embedding on the other ball. The 3-handle can be attached as follows. The caps D_{\pm}^2 of the attaching sphere are in the drawing canvas of the 0-handle, bounding S^1 on D_+^3 and its mirror circle on D_-^3 , respectively. The cylinder $S^1 \times [+1, -1]$ vanishes inside the remaining region of the 1-handle, conforming to the Kirby conventions.

3-2 It is helpful to visualise the 3-2-slide by realising the attaching sphere of h_3 as the “double pancake” $S^2 \cong D^2 \cup_{S^1} D^2$, which consists of two discs glued along their boundary. The gluing S^1 is then slid over the attaching S^1 of the 2-handle, as in the 2-2-move, dragging the discs along. This is illustrated in Figure 10.

For an explicit model, assume that the 2-handle h_2 is attached to the handlebody H along $\partial_a h_2 = S^1 \times D^2$. Choose an embedded interval $[-1, +1] \subset \partial D^2$. This defines a strip $S^1 \times [-1, +1]$ on the boundary of $\partial_a h_2$, and a thickened disc $D^2 \times [-1, +1]$ inside the remaining region $\partial_r h_2$. The 3-handle now isotopes from the old attaching sphere S_0^2 to the new attaching sphere S_n^2 , and we seek the image of this isotopy in the drawing canvas, which is a ball with a thickened disc and a ball removed:

$$M := (D_n^3 \setminus D^2 \times [-1, +1]) \setminus D_0^3.$$

The $D^2 \times [-1, +1]$ is understood to be embedded such that the end discs $D^2 \times \{-1, +1\}$ are embedded in the boundary of D_n^3 , one on each “pancake.” M has as boundary two components, $(S_n^2 \setminus (D_+^2 \sqcup D_-^2)) \cup S^1 \times [-1, +1]$, and S_0^2 . Assuming that the 3-handle is initially attached along S_0^2 , it can be isotoped such that it is partially in the remaining region of the 2-handle, and the part staying visible in the main drawing canvas is $S_n^2 \setminus (D_+^2 \sqcup D_-^2)$. It vanishes into $\partial_r h_2$ at S^1 times the endpoints of $[-1, +1]$.

3-3 One attaching S^2 of a 3-handle can “engulf” another one. In analogy to the 2-2-slide, the three-dimensional pair of pants is defined as a ball D_C^3 with two open balls $D_A^{2,\circ}$ and $D_B^{2,\circ}$ removed. Its boundary is $S_A^2 \sqcup S_B^2 \sqcup S_C^2$. Again, an embedding of the three-dimensional pair of pants in the boundary of a handlebody defines three 3-handle attachments $h_{3,A}$, $h_{3,B}$ and $h_{3,C}$. The handle $h_{3,A}$ can be slid over $h_{3,B}$, which replaces it by $h_{3,C}$. This is illustrated in Figure 11.




Figure 10. The 3-2-slide. Graphically, we can imagine the fold arc of the 3-handle attachment slides over the 2-handle. The 3-handle is understood to be extended past the dashed lines, to form a whole sphere.




Figure 11. In the 3-3-slide, one 3-handle attachment can slide over the other.




Figure 12. Two Kirby diagrams arising from the same handle decomposition of $S^1 \times S^3$, related by the 3- ∞ -move.

3- ∞ Regularity requires handle attachments not to intersect with the point at infinity of $\partial_r h_0 = S^3 \cong \mathbb{R}^3 \cup \{\infty\}$. While any isotopy of 1-handles and 2-handles can be deformed away from this point such that the attachment never intersects with it, this is not always possible for 3-handles. The fundamental move that isotopes the attaching S^2 over this point is called the 3- ∞ -move.

When a 3-handle is only attached to h_0 , its attaching S^2 separates the drawing canvas into an inside region and an outside region. The 3- ∞ -move interchanges the roles of the two regions, as illustrated in Figure 12.

Remark 3.13. The 3- ∞ -move is technically not a handle slide, but shares enough similarity in order to list it with the other slides. It is helpful to visualise the move applied to a 3-handle h_3 as the following procedure:

- (1) attach a new 3-handle h_3^∞ such that its attaching S^2 surrounds the whole diagram, or equivalently the point at ∞ ;
- (2) slide h_3 over h_3^∞ ;
- (3) remove h_3^∞ .

Remark 3.14. Often it is possible to perform slides along existing handles. For example, sliding a 1-handle $h_{1,\alpha}$ over $h_{1,\beta}$ requires a path from $\partial_a h_{1,\alpha}$ to $\partial_a h_{1,\beta}$. If there is a 2-handle h_2 attached to both these 1-handles, a segment of its attaching circle provides such a path, and the 1-1-slide can be performed along it, while simultaneously sliding h_2 over $h_{1,\beta}$ as well.

Similarly, a 2-2-slide can sometimes be performed along a 3-handle attachment, while simultaneously sliding the 3-handle. This is illustrated in Figure 9.

3.4. Examples

3.4.1. $S^1 \times S^3$. There is a standard handle decomposition of S^n as a single 0-handle and a single n -handle. Since handle decompositions have products, and $h_{k_1} \times h_{k_2} \cong h_{k_1+k_2}$, there is a decomposition of $S^1 \times S^3$ into $h_0 \cup h_1 \cup h_3 \cup h_4$. Two possible Kirby diagrams of this handle body are shown in Figure 12. They are related via the 3- ∞ -move. See also [15, Figure 4.15].

3.4.2. $S^1 \times S^1 \times S^2$. For 3-dimensional manifolds, Heegard diagrams are a standard depiction of handle decompositions. [15, Example 4.6.8] explains how to construct a handle decomposition of $S^1 \times M^3$ from a Heegard diagram of M^3 . In [5, Section 6.2.1], a Kirby diagram of $S^1 \times S^1 \times S^2$ is derived, but 3-handles are suppressed. Figure 13 shows the same Kirby diagram with 3-handles.

The reader might be surprised about the apparent lack of symmetry between D_α^{3+} and D_α^{3-} (and similarly D_β^{3+} and D_β^{3-}), with the 3-handle spheres covering only one of the 1-handle balls. (After all, each 1-handle corresponds to an S^1 factor, and mirroring the diagram such that, e.g., $D_\alpha^{3\pm}$ are preserved and the $D_\beta^{3\pm}$ are interchanged results in orientation reversal of the S^1 corresponding to the 1-handle $h_{1\beta}$, which is an isomorphism.) This asymmetry stems from the second regularity condition in Definition 3.7, which stipulates that no handle attachment intersect with ∞ . If we remove that restriction, we can slide a point on $\partial_a h_{2b}$ to ∞ and end up with a symmetric diagram like the one given in [16, Figure 4.1a], where h_{2b} is attached along the y -axis, and h_{3A} and h_{3B} are attached along the $y - z$ -plane and the $x - y$ -plane, respectively.

3.4.3. Fundamental group. It is well known that the 2-handlebody of a handle decomposition gives a presentation of the fundamental group [15, Solution of Exercise 4.6.4(b)], [5, Section 2.3.3], where each 1-handle is a generator and each 2-handle a relation.

For example, it is a good exercise to verify that in Figure 13, the two 1-handles $h_{1\alpha}$ and $h_{1\beta}$ constitute two generators α and β each, and the 2-handle h_{2a} results in the relation $\alpha\beta\alpha^{-1}\beta^{-1}$ (after having oriented the attaching spheres arbitrarily). h_{2b} is not attached to any 1-handles and thus yields the trivial relation. We have shown that $\pi_1(S^1 \times S^1 \times S^2) \cong \mathbb{Z} \oplus \mathbb{Z}$, as expected.

Turning a handle decomposition upside down [15, Section 4.2] shows that it is possible to present the fundamental group with the 3-handles as generators and again the 2-handles as relations. Each 2-handle yields a relation where the generators corresponding to all attached 3-handles are multiplied in cyclical order.⁵ A generator is inverted if the boundary orientation deriving from the (arbitrarily chosen) orientation of the attaching sphere S^2 of h_3 does not match the (again arbitrarily chosen) orientation of the attaching S^1 of h_2 .




Figure 13. Handle decomposition of $S^1 \times S^1 \times S^2$ with two 3-handles, two 2-handles and two 1-handles. Choosing an arbitrary orientation and start point, the attaching S^1 of h_{2a} attaches to $h_{1\alpha}$ at D_α^{3-} , runs through the remaining region $\partial_r h_{1\alpha}$, leaves at D_α^{3+} , attaches to $h_{1\beta}$ at D_β^{3+} , runs through $\partial_r h_{1\beta}$, leaves at D_β^{3-} , attaching a second time to $h_{1\alpha}$ and $h_{1\beta}$ each, and finally closing the loop. The attachment of h_{3A} starts (invisibly) inside the remaining region of h_{2b} , extends to a D^2 and leaves it at $\partial_a h_{2b}$, proceeding as a cylinder around $\partial_a h_{2a}$ and D_α^{3+} , attaching to $h_{1\beta}$ at D_β^{3+} , leaving again at D_β^{3-} and attaching a second time to h_{2b} , closing off inside its remaining region.

⁵To see this, visualise that for each 3-handle, a noncontractible S^1 starts in the interior of the single 4-handle, passes through the first of the two D^3 s constituting the remaining region of the 3-handle, and becomes visible as $\text{pt.} \times [-1, +1] \subset S^2 \times [-1, +1]$ in the thickened attaching sphere of the 3-handle, and enters the second D^3 of the remaining region, back into the 4-handle. It can be contracted in the remaining region of a 2-handle if the 3-handle is attached to it.

Again, Figure 13 offers a good exercise to compute the fundamental group in this presentation. The two 3-handles yield generators A and B . None of the 3-handles is attached to h_{2a} , so its relation is trivial. Both 3-handles are attached twice to h_{2b} , though. Choosing arbitrary orientations and taking care that the orientations on both parts of the attaching spheres of the 3-handles match, we read off the relation $ABA^{-1}B^{-1}$ and arrive at the same group.

4. Graphical calculus in G -crossed braided spherical fusion categories

Our strategy to define an invariant of 4-manifolds from a G -crossed braided spherical fusion category \mathcal{C} (short: $G\times$ -BSFC) is to choose a Kirby diagram of the manifold, and to interpret and evaluate this diagram in a graphical calculus of \mathcal{C} . As will be shown in this section, this calculus is conveniently similar to Kirby calculus with 3-handles.

4.1. From manifolds to morphisms

Our general guiding principle to arrive at a graphical calculus is that $G\times$ -BSFCs are special degenerate (weak) 3-categories, and thus monoidal 2-categories [11]. Viewing them as 3-categories, there is a single object, the 1-morphisms correspond to group elements, the 2-morphisms correspond to objects of the fusion category, and the 3-morphisms correspond to morphisms of the fusion category. The diagrams should thus be 3-dimensional, with group elements labelling 2-dimensional *sheets*, objects of the fusion category labelling 1-dimensional *ribbons*, and morphisms of the fusion category labelling *points*.

Here, a similarity can already be seen to the elements of Kirby diagrams, which are the S^2 s corresponding to 3-handles, 1-dimension curves for 2-handles, and D^3 s (thick points) for 1-handles. Kirby diagrams themselves are not flexible enough for arbitrary graphical calculus in a $G\times$ -BSFC. This leads us to the definition of planar diagrams, which are an intermediate step between Kirby diagrams and the desired graphical calculus. Figure 14 gives an overview over all steps necessary to compute the invariant from a manifold, which are an intermediate step between Kirby diagrams and the desired graphical calculus.

4.2. Unlabelled planar diagrams

Definition 4.1. An unlabelled, closed *diagram* \mathcal{D} (see Figure 15) consists of the following data:

- (0) a finite set \mathcal{D}_0 of embeddings $p: D^3 \rightarrow \mathbb{R}^3$, called (thickened) *points*,

- (1) a finite set \mathcal{D}_1 of framed embeddings of *lines* $[-1, 1]$ and *circles* S^1 (collectively called *ribbons*) into \mathbb{R}^3 such that the boundaries (endpoints) of each line are embedded in the boundaries of the thickened points,
- (2) a finite set \mathcal{D}_2 of smooth embeddings of n -gons (see Definition A.1), called *sheets*, or sometimes *discs* (when $n > 0$) and *spheres* into \mathbb{R}^3 such that each boundary component of each disc is either a line or embedded on a boundary of a thickened point.

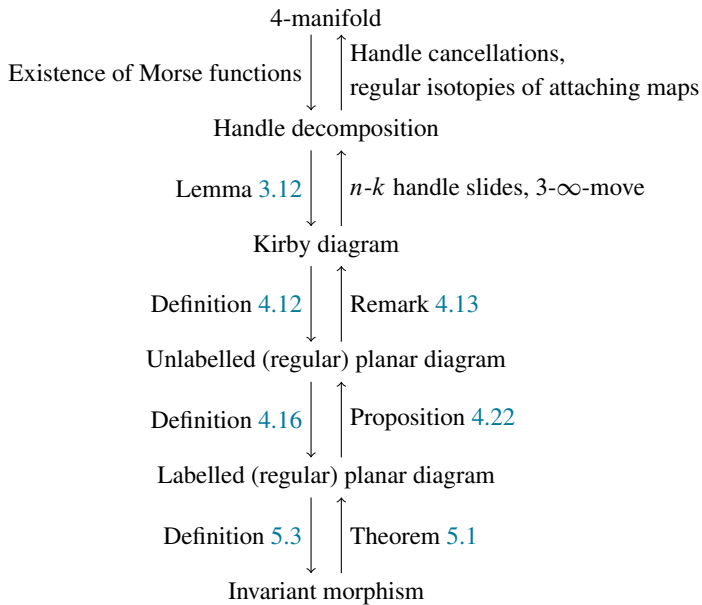


Figure 14. In order to define the invariant, a 4-manifold is translated to a morphism in a G -BSFC through various intermediate steps. For each of these steps, additional data has to be added, or additional properties assumed. The arrows pointing down illustrate this additional structure, while the arrows pointing up justify it through a proof of independence on the specific structure chosen.



Figure 15. An unlabelled closed diagram with a single point, a single line, and a single disc.

Remark 4.2. n -gons are manifolds with corners, likewise their embedded image will have corners.⁶ The n -gons will always be $2k$ -gons, in fact, where k is the number of times the sheet attaches to a thickened point (or equivalently the number of lines it attaches to). The corners of the sheet will always be attached to an endpoint of a line embedding.

The reader may be concerned at this point that neither sheets nor ribbons were required to be thickened, or framed. The reason is that discs and spheres contain no framing information, and the framing of a ribbon can be canonically chosen to be the *blackboard framing* by choosing a projection into the plane. (For further details, also see Appendix B.2.) Since we will ultimately want to evaluate diagrams in monoidal categories (which have a diagrammatic calculus in the plane), we shall make sure that the projection of a diagram contains all essential information. This is encoded in the following definition.

Definition 4.3 (Blackboard framing). We single out a projection into the plane onto the first two components of a vector:

$$\begin{aligned} \pi: \mathbb{R}^3 &\xrightarrow{\cong} \mathbb{R}^2 \oplus \mathbb{R} \xrightarrow{\pi_{\mathbb{R}^2}} \mathbb{R}^2, \\ (x, y, z) &\mapsto (x, y). \end{aligned}$$

A diagram \mathcal{D} is compatible with the *blackboard framing* if the following conditions are satisfied.

- Each thickened point D^3 can be split in half along an *equator* S^1 such that its boundary is split into “upper” and “lower” discs, $S^2 = D_u^2 \cup_{S^1} D_l^2$ and each disc $D_{\{u,l\}}^2$ is *embedded* into \mathbb{R}^2 by π .
- Lines attach to thickened points transversely in the equator.
- On sheets, π is an immersion except at a finite set of 1-dimensional compact embedded piecewise smooth manifolds, the *fold graphic* [27, Section 1.3]. The non-smooth points of the fold graphic are called *cusps points*, and the smooth parts *fold arcs*.
- On the union of fold arcs and ribbons, π is an embedding except at a finite set of points, the *crossings*.

Remark 4.4. In the projection, a thickened point is given by simply an embedded D^2 .

⁶Instead of smooth embeddings of manifolds with corners, one may imagine a *piecewise smooth* embedding of a disc or sphere instead. The formulation with corners appears cleaner though, particularly in the light of Definition 3.1, where handles are defined as manifolds with corners.

Remark 4.5. The embedded diagram data are understood to be oriented. The resulting data is sometimes depicted graphically, which manifests depending on the dimension on the datum:

- (0) the orientation of points can either agree with the orientation of \mathbb{R}^3 , or be opposite, specifying a sign;
- (1) the orientation on a ribbon manifests as a direction along the ribbon, often depicted with a little arrow;
- (2) the orientation of a sheet can coincide locally with the orientation of the plane \mathbb{R}^2 , or be opposite, specifying a sign per noncritically embedded stratum.

Remark 4.6. Ribbons (and borders of thickened points) locally partition the plane into two “sides.” π being an immersion on sheets outside of fold arcs implies that sheets incident to ribbons never change the side locally.

This may be confusing at first when recalling that we emphasised in Remark 3.10 that the single picture conventions (Definition 3.9) imply that in a 3-2-handle attachment, the 3-handle must follow the framing of the 2-handle. But since any ribbon also follows the blackboard framing, it does not “turn” in the plane, so these two conditions conform.

The parts where π is critical on the diagram, the fold graphic and the crossings, will later correspond to structure morphisms of the G -crossed fusion category. Certain regularity conditions need to be imposed in order to make the diagrams suitable for graphical calculus.

Definition 4.7 (Regular diagrams). A diagram \mathcal{D} is *regular* if the following conditions are satisfied:

- it is compatible with the blackboard framing;
- the projections of ribbons and fold arcs do not intersect projections of thickened points;
- the projections of exactly two ribbons, two fold arcs, or a ribbon and a fold arc, intersect transversely at a crossing;
- the projections of different crossings and cusp points never coincide.

Remark 4.8. One can convince oneself that one can perturb any planar diagram with an isotopy to a regular diagram. They are thus also often said to be in *general position*, and we will usually assume that a diagram is regular.

Definition 4.9. A *regular planar isotopy* is an isotopy of planar diagrams, such that they are regular, and each intermediate diagram is regular.

Lemma 4.10. *Two regular planar diagrams are framed isotopic if and only if they are related by a finite sequence of regular planar isotopies and planar moves, illustrated in Figures 16 and 17.*

Proof. For points and ribbons, this is standard, e.g., [17, Theorem 3.7] [28, Theorem 6.1]. For sheets, this is an exercise in Cerf theory, carried out in detail, e.g., in [27, Chapter 1.5]. ■

Remark 4.11. The Reidemeister I move, which equates a single twist by 2π of one ribbon to an untwisted ribbon, does not appear among the list of moves. This is intentional. While it is important for unframed knot and link invariants, it is not admissible in framed invariants since it changes the framing. The manifold invariant developed here is sensitive to the framing (otherwise it would not be able to distinguish $\mathbb{C}\mathbb{P}^2$ and S^4), and thus cannot validate the Reidemeister I move.

In ribbon categories, this is mirrored algebraically by the fact that in the non-symmetric case, the twist is not trivial. In $G \times$ -BSFCs, the twist is not even an endomorphism canonically. For a g -graded object $X \in \mathcal{C}_g$, it types as $\theta_X: X \rightarrow {}^g X$, and there is in general no coherence isomorphism to compare it to. (While X and ${}^g X$ are isomorphic, they are not canonically isomorphic if g is nontrivial.)

We have drawn diagrams of Kirby diagrams before in Section 3. They are indeed diagrams in the sense of Definition 4.1.

Definition 4.12. A Kirby diagram K (in the sense of Definition 3.11) of M straightforwardly gives rise to an unlabelled planar diagram \mathcal{D} , the *derived* planar diagram. It is computed by translating the attaching spheres to elements of the diagram.

- (1) Every 1-handle h_1 gives rise to two thickened points D_{\pm}^3 , for each ball of the attaching region of h_1 .
- (2) A 2-handle gives rise to an embedded circle (a knot) if it is not attached to any 1-handles, or else to one or several lines incident on the thickened points corresponding to the 1-handles it is attached to. Possibly, an isotopy has to be applied such that the blackboard framing of the embedding matches the given framing of the attachment.
- (3) A 3-handle gives rise to an embedded sphere if it is not attached to any handles, or else to one or several discs incident on the thickened points and ribbons corresponding to the 1-handles and 2-handles it is attached to. Possibly, an isotopy has to be applied such that in the projection, no disc “changes the side” of a 2-handle it is attached to (see Appendix B for details).

Remark 4.13. Isotopies of Kirby diagrams correspond to isotopies of planar diagrams, so Lemma 4.10 again applies: two Kirby diagrams are isotopic if they are related by regular isotopies and planar moves.

4.3. Labelled diagrams

If a regular diagram is labelled appropriately with data from a $G \times$ -BSFC \mathcal{C} , a morphism in \mathcal{C} can be extracted. This situation is set up in the following two definitions.

Definition 4.14. A sheet with boundary on a ribbon r is called *incident* to r . All sheets $\{(s_i, \pm)\}$ incident to r are written as δr . They form an ordered multiset, starting at the top (viewed from the projection) right-hand (viewed from the ribbon orientation) sheet and going completely around the ribbon with the right-hand rule (starting, at first, into the drawing plane). The sign is $+$ if the boundary orientation of the sheet matches the orientation of the ribbon, and $-$ otherwise.

Analogously, the ribbons $\{(r_i, \pm)\}$ incident to a thickened point p are denoted as δp . This is a cyclically ordered set, starting anywhere on the boundary S^1 of the projected disc of p and proceeding counter-clockwise. The sign encodes with which endpoint the ribbon attaches to the point. (It may attach with both ends.)

Definition 4.15. A *labelling* of a diagram \mathcal{D} with a G -crossed fusion category \mathcal{C} consists of three functions with the following signatures:

$$g: \mathcal{D}_3 \rightarrow G, \quad X: \mathcal{D}_2 \rightarrow \mathcal{O}(\mathcal{C}), \quad \iota: \mathcal{D}_1 \rightarrow \text{mor } \mathcal{C}.$$

They need to satisfy the following *typing relations*:

$$\deg(X(r)) = \prod_{(s, \pm) \in \delta r} g(s)^\pm, \quad \iota(p) \in \left\langle \bigotimes_{(r, \pm) \in \delta p} X(r)^\pm \right\rangle.$$

Here, $g(s)^\pm$ denotes either $g(s)$ or $g(s)^{-1}$, depending on the sign of (s, \pm) . Similarly, by $X(r)^\pm$ we mean $X(r)$ for $+$ and $X(r)^*$ for $-$. The product is performed in the order specified in the previous definition.

Informally, a labelling attaches group elements to sheets, simple objects to ribbons, and morphisms to the points. The degree of the objects on a ribbon is given by the sheets incident to the ribbon, and the morphism of a point must be in the morphism space of the objects labelling the incident ribbons.

To extract a morphism from a labelled diagram, we project the diagram into the plane and mostly apply the well-known diagrammatic calculus of pivotal fusion categories, treating the crossings and sheets as additional data. These translate directly to G -crossed structures.













| Name | Left-hand side | Right-hand side |
|---|---|--|
| Point isotopy |  |  |
| Reidemeister I' (twist move) |  |  |
| Reidemeister II (cancellation of inverse crossings) |  |  |
| Reidemeister III (crossing isotopy) |  |  |
| Cusp isotopy |  |  |
| Cusp moves. Shown here: cancellation/inversion. Also relevant: Other inversion side, flip, swallowtail. |  |  |

Figure 16. Different kinds of planar moves. Note that on each ribbon, there are infinitely many possibilities of attaching further 3-handles, for which only one example per move is shown. Furthermore, each move that occurs for a ribbon can also occur for a fold line.

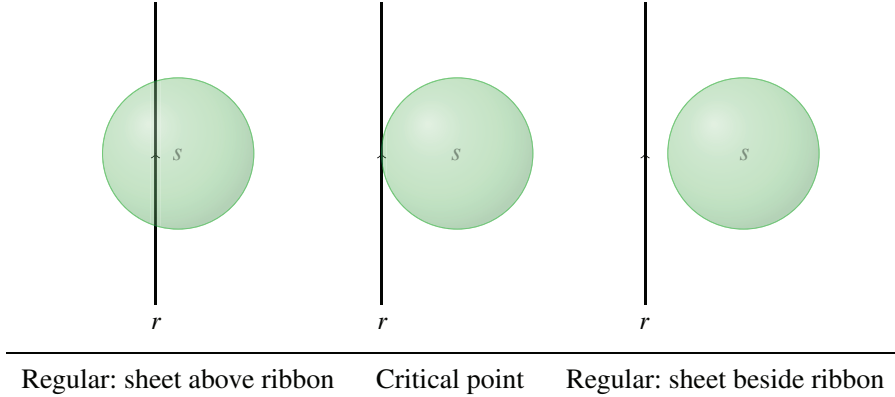


Figure 17. By changing the projection, a regular diagram can change into one that does not satisfy the regularity condition. In this case, the intersection of the projection of a fold arc and a ribbon is not transverse. A further change leads again to a regular diagram, but one that is not regularly isotopic to the original one. They are related by a planar move though, in this case the Reidemeister II move.

Definition 4.16. A labelled, regular diagram \mathcal{D} can be evaluated to an endomorphism of \mathcal{I} (a complex number). To evaluate the diagram for given g , X and ι , follow this algorithm.

- (1) Starting from the back of the drawing plane, whenever a sheet s covers ribbons or points, $g(s)^\pm$ acts on objects and morphisms labelling them. (The sign specifies whether the projection maps the sheet orientation onto the canonical orientation of the plane.)
- (2) Insert a crossed braiding for an intersection point of two ribbons (Figure 18), and appropriate G -crossed coherence isomorphism for every sheet incident on the right-hand side of the overcrossing ribbon. (Figure 19).
- (3) Insert appropriate G -crossed coherence isomorphisms at intersections involving at least one sheet fold arc. (Figure 20). (For cusps, no morphism has to be inserted.)
- (4) Interpret the resulting diagram in the graphical calculus of pivotal fusion categories.

The resulting endomorphism of \mathcal{I} , and equivalently the corresponding number, is denoted as $\langle \mathcal{D}(g, X, \iota) \rangle$.

Remark 4.17. Without difficulty, this definition can be generalised to open diagrams as in Section 2.1. The result is then in a morphism space $\langle A \otimes B \otimes \dots \rangle$, where A, B, \dots , are the open-ended ribbons in the diagram.




Figure 18. Evaluating the crossed braiding for the sheet labelling g and the ribbon labelling X . The typing constraints demand that $\deg(X(r_1)) = g(s_1)$ and $\deg(X(r_2)) = g(s_2)$.




Figure 19. For sheets on the right-hand side of an overcrossing ribbon, an additional coherence has to be inserted. Note that here, $\deg(X(r_1)) = g(s_1)^{-1}$, as becomes apparent when considering a coordinate patch in which s_1 and r_1 are mapped on the upper half plane and the x -axis.

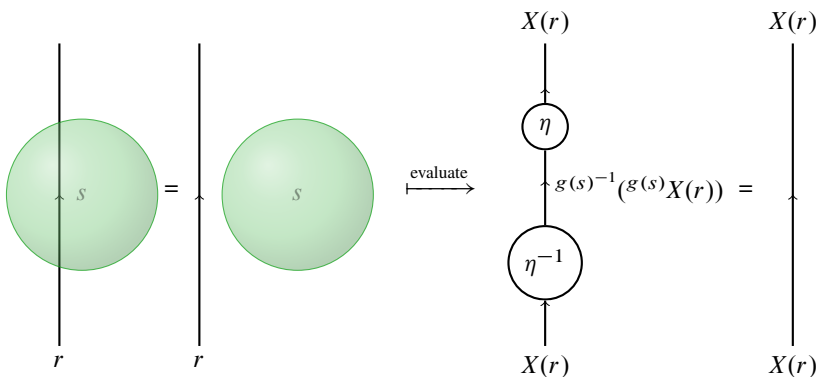


Figure 20. An $s = S^2$ sheet covers the ribbon r . The covered interval of the ribbon is acted upon twice, with $g(s)^{-1}$ and $g(s)$. (Uniquely determined) G -coherence morphisms are inserted at the crossings of fold arc and ribbon. Once isotoped away, s does not cover any other part of the diagram and thus acts as the identity.




Figure 21. The twist is an isomorphism, but not canonically an automorphism. This can be seen from the different covering of sheets on the ends of the ribbon: when the ribbon performs a full twist, the sheet wraps it once fully. This adds a group action on the target object.

Remark 4.18. Since $G \times$ -BSFCs satisfy a coherence theorem ([24], see also [11, Theorem 2.3]), the choice of coherence isomorphism is always uniquely defined, given all group labellings g and object labellings X .

Remark 4.19. A sheet incident on a point does not act on the morphism that labels the point, but only on the objects (resp. morphisms) labelling the ribbons (resp. points) which are *covered* by the sheet in the projection.

Remark 4.20. Fold lines play a surprisingly small role in this calculus. It is usually helpful to imagine fold lines to be labelled with the monoidal unit I of the category. This makes it more intuitive why they do not influence the evaluation.

Examples are found in Section 6.

Lemma 4.21. *The evaluation of a labelled diagram is invariant under isotopies of regular diagrams (Definition 4.9).*

Proof. Regular isotopies do not change the topology of the projected diagram in the plane. By coherence of the diagrammatic calculus of pivotal fusion categories, the evaluation is invariant. ■

Proposition 4.22. *The evaluation of a diagram is invariant under arbitrary (framed) isotopies of diagrams.*

Proof. By Lemma 4.10 and the previous lemma, we need to prove invariance of the diagram evaluation from Definition 4.16 under each planar move.

Point isotopy. Naturality of the crossed braiding.

Reidemeister I'. By definition of the twist, see, e.g., [11, Figure 4a] for a definition and Figure 21 for its evaluation.

Reidemeister II. By the definition of the graphical representation of the inverse braiding (in the case of ribbons), or the inverse of the coherence η , as seen in Figure 20.

Reidemeister III. By the G -crossed braid axiom, sometimes also called the “heptagon axiom,” as in [11, Definition 2.2.2].

Cusp isotopy. The left-hand side of this move evaluates to two cancelling coherence isomorphisms.

Cusp cancellation. When the folded sheet parts do not cover any other part of the diagram (as can be achieved with a combination of other moves such as cusp isotopies), then these parts do not change the diagram. ■

Remark 4.23. Once a sheet labelling g and a ribbon labelling X are fixed, a labelling ι of points with morphisms gives rise to an elementary tensor in the tensor product of all morphism spaces in the diagram

$$\bigotimes_{p \in \mathcal{D}_1} \iota(p) \in \bigotimes_{p \in \mathcal{D}_1} \left(\bigotimes_{(r, \pm) \in \delta p} X(r)^\pm \right).$$

It will turn out that arbitrary vectors in the right-hand side vector space are often relevant, so we will generalise the notion of morphism labellings to arbitrary vectors, not just elementary tensors. The evaluation from the previous definition can be generalised uniquely to such vectors, by the universal property of the tensor product.

Similarly, object labellings X are generalised to a labelling with elements from the fusion algebra $\mathbb{C}[\mathcal{C}]$. From now on, we will often implicitly make use of these two generalisations.

4.4. Kirby colours and sliding lemmas

Let us fix some notation and prove the essential lemmas for the invariant definition in Section 5. For this section, assume \mathcal{C} to be a $G \times$ -BSFC.

4.4.1. G -graded fusion categories. Recall that the set of (chosen representatives of equivalence classes of) simple objects in \mathcal{C} is denoted by $\mathcal{O}(\mathcal{C})$, and the same notation applies to any semisimple linear category. Set cardinality is denoted as $|\mathcal{O}(\mathcal{C})|$.

Definition 4.24 ([29, Definition 6.6]). The *Kirby colour* of degree g is defined as the following element of the fusion algebra $\mathbb{C}[\mathcal{C}]$:

$$\Omega_g := \Omega_{\mathcal{C}_g} = \bigoplus_{X \in \mathcal{O}(\mathcal{C}_g)} d(X)X.$$

Remark 4.25. Note that $\Omega_{\mathcal{C}} = \bigoplus_{g \in G} \Omega_g$.

Remark 4.26. Unlike $\Omega_{\mathcal{C}}$, Ω_g is not self-dual. In fact, sphericity implies $\Omega_g^* \cong \Omega_{g^{-1}}$. This implies that in the graphical calculus, the orientation of ribbons labelled Ω_g needs to be specified.

Lemma 4.27. *Let $\mathcal{C}_g \neq 0$. Then $d(\Omega_g) = d(\Omega_e)$.*

Proof. Jumping slightly ahead and using Lemma 4.31, we slide a loop labelled with Ω_g over an Ω_e -loop and find $d(\Omega_g)^2 = d(\Omega_g)d(\Omega_e)$. Then $d(\Omega_g)$ can be cancelled since it is nonzero. For the analogous lemma with Frobenius–Perron dimensions, see [14, Proposition 8.20]. ■

Remark 4.28. For a faithful grading ($\mathcal{C}_g \neq 0$ for every $g \in G$), the previous lemma is equivalent to $d(\Omega_g) \cdot |G| = d(\Omega_{\mathcal{C}})$. But since a non-faithful G -grading on \mathcal{C} is always given by a subgroup of $H \subset G$ and a faithful H -grading, we can still leverage the equation $d(\Omega_g) \cdot |H| = d(\Omega_e)$ if this subgroup is known.

4.4.2. The G -crossed braiding and encirclings.

Definition 4.29. Let $h, g \in G$. The double braiding of two objects $A \in \mathcal{C}_h$ and $B \in \mathcal{C}_g$ is defined as

$$\beta_{A,B}: A \otimes B \rightarrow {}^{hg}A \otimes {}^hB, \quad \beta_{A,B} := (\eta_B \otimes 1_B) \circ c_{B,A} \circ c_{A,B}.$$

η_B is the unique coherence of the G -action. Note that $\beta_{A,B}$ is, up to $G \times$ -coherences, an automorphism if $h = g = e$.

Definition 4.30. Let $g \in G$. The *encircling* of an object $A \in \mathcal{C}_e$ by an object $B \in \mathcal{C}_g$ is defined by the following partial trace:

$$\Delta_{A,B}: A \rightarrow {}^gA, \quad \Delta_{A,B} := \text{tr}_B((1_A \otimes \varepsilon_B) \circ \beta_{A,B}): A \rightarrow {}^gA.$$

B can be generalised to (g -graded) elements of the fusion algebra straightforwardly, and we will freely make use of this generalisation.

Lemma 4.31. *Assume $A \in \mathcal{C}_e$, $h, g \in G$ and $B \in \mathcal{C}_g$. There is a graded sliding lemma for $G \times$ -BSFCs, which changes the grade of the encircling:*

$$\Delta_{A,\Omega_h} \otimes 1_B = (\eta_A \otimes 1_B) \circ c_{B,A} \circ (\varepsilon_B \otimes \Delta_{A,\Omega_{g^{-1}h}}) \circ c_{A,B}. \quad (4.1)$$

The encirclement may be arbitrarily linked or knotted, as far as this is possible with regard to the grading.

Proof. In analogy to usual proofs of sliding lemmas (e.g., [22, Corollary 3.5] and [5, Lemmas 3.3 and 3.4]), B^* and Ω_h are fused to a single strand, but since $B \in \text{ob } \mathcal{C}_g$, this strand has to be labelled $\Omega_{g^{-1}h}$. We have used [29, Lemma 6.6.1]. ■

Remark 4.32. To see why the previous lemma is called “sliding lemma,” revisit Figure 9, which contains graphical representations of the two sides of (4.1) if the 3-handle h_{3A} is labelled with $g^{-1}h$, h_{3B} with h , h_{2a} with B , and the gray area is replaced by a ribbon labelled with A .

Definition 4.33 (Well known, e.g., [5, Definition 2.41]). Let \mathcal{D} be a braided fusion category. Then \mathcal{D}' is the full symmetric subcategory spanned by trivially braiding objects, called the *symmetric centre*.

Let A be an object in \mathcal{D} . Then A' is defined (up to isomorphism) to be the maximal subobject of A in \mathcal{D}' , and $\tau_A: A \rightarrow A'$ the idempotent defined as projection onto A' followed by inclusion into A (not depending on the choice of A'). If $\tau_A = 1_A$, or equivalently $A \cong A'$, or $A \in \text{ob } \mathcal{D}'$, then A is said to be *transparent*.

Lemma 4.34 (Well known, e.g., [5, Lemma 2.46]). *Let \mathcal{D} be a ribbon fusion category and A an object therein. Then $\Delta_{A, \Omega_{\mathcal{D}}} = \tau_A \cdot d(\Omega_{\mathcal{D}})$. In particular, let X be a simple object in \mathcal{D} . Then $\Delta_{X, \Omega_{\mathcal{D}}} = 1_X \cdot d(\Omega_{\mathcal{D}})$ if and only if X is transparent, and 0 otherwise.*

This lemma is called the killing lemma, since nontransparent X are “killed off” by an encircling with $\Omega_{\mathcal{D}}$.

Lemma 4.35. *To the knowledge of the author, this generalisation of the killing lemma has not been discussed in the literature before.*

Assume that $g \in G$ and $\mathcal{C}_g \neq 0$. Then the $G \times$ -killing lemma holds for any $A \in \text{ob } \mathcal{C}_e$:

$$\Delta_{A, \Omega_{g^{-1}}} \circ \Delta_{A, \Omega_g} = \tau_A \cdot d(\Omega_e)^2.$$

Proof. The left-hand side is represented diagrammatically by a closed cylinder encircling an A -labelled ribbon, with the sheet labelled g . The $\Omega_{g^{-1}}$ -encircling can be slid off the Ω_g -encircling, yielding factor $d(\Omega_{g^{-1}}) = d(\Omega_g)$. This changes the grade of the remaining encircling to Ω_e , thus we can apply the killing lemma (Lemma 4.34). ■

Corollary 4.36. *Let $X \in \mathcal{C}_e$ be simple. Then encircling Δ_{X, Ω_g} is an isomorphism $X \rightarrow {}^g X$ if and only if X is transparent (the inverse being $\Delta_{X, \Omega_{g^{-1}}} \cdot d(\Omega_e)^{-2}$), and 0 otherwise. This justifies generalising the name “killing lemma.”*

5. The invariant

In this section, we will define an invariant of closed, smooth, oriented 4-manifolds by labelling Kirby diagrams (Section 3) with data from $G \times$ -BSFCs (Sections 2.2

and 4.3), and show its independence of the chosen Kirby diagram by means of the lemmas from Section 4.4.

Fix again a $G \times$ -BSFC \mathcal{C} , a manifold M , and a Kirby diagram K for M . Recall from Definition 4.12 that K gives rise to an unlabelled planar diagram \mathcal{D} . The 1-, 2-, and 3-handles of K can be labelled appropriately to yield a labelling of the derived diagram \mathcal{D} . In detail, we require the following data.

Definition 5.1 (Labelling of Kirby diagrams). Denote by K_j the set of j -handles. A labelling of K by \mathcal{C} is specified by

- a function $g: K_3 \rightarrow G$,
- a function $X: K_2 \rightarrow \mathbb{C}[\mathcal{C}]$ such that the derived labelling of \mathcal{D} type-checks,
- a function $\iota: K_1 \rightarrow \text{mor } \mathcal{C} \otimes \text{mor } \mathcal{C}$ into type-checking dual morphism spaces:

$$\iota(h_1) \in \langle \otimes_{(h_2, \pm) \delta h_1} X(h_2)^{\pm} \rangle \otimes \langle \otimes_{(h_2, \pm) \delta^{-1} h_1} X(h_2)^{\mp} \rangle.$$

Here, δh_1 denote the attached 2-handles, and δ^{-1} denotes the reverse cyclical ordering.

Remark 5.2. Crucially, the handles of the Kirby diagrams receive labels, and not each individual element of the diagram. It is important that the different lines stemming from a single 2-handle are labelled with the same data, and the same is true for the different discs from a single 3-handle.

The invariant is now defined as a sum of diagram evaluations over all possible 3-handle labellings, assigning the appropriately graded Kirby colour to every 2-handle and the dual bases from Definition 2.1 to every 1-handle.

Definition 5.3. Making full use of the generalisations from Remark 4.23, the invariant assigned to a $G \times$ -BSFC \mathcal{C} is defined as

$$I_{\mathcal{C}}(K) := \sum_{g: K_3 \rightarrow G} \left\langle K \left(g, h_2 \mapsto \Omega_{\deg(h_2)}, \bigotimes_{h_1} \sum_i \phi_{h_1, i} \otimes \tilde{\phi}_{h_1, i} \right) \right\rangle d(\Omega_{\mathcal{C}})^{|K_1| - |K_2|}.$$

By abuse of notation, we have used $\deg(h_2) := \prod_{(g, \pm) \in \delta h_2} g^{\pm}$, which is the degree of a type-checking object labelling a 2-handle h_2 . $\sum_i \phi_{h_1, i} \otimes \tilde{\phi}_{h_1, i}$ denotes the sum over the dual bases of $\langle s(h_1) \rangle$ and $\langle s(h_1)^* \rangle$. Writing out the sums explicitly is possible:

$$\begin{aligned} I_{\mathcal{C}}(K) = & \sum_{g: K_3 \rightarrow G} \sum_{\substack{X: K_2 \rightarrow \mathcal{O}(\mathcal{C}) \\ \deg(X(h_2)) = \deg(h_2)}} \sum_{\substack{\phi_{h_1, i}, \tilde{\phi}_{h_1, i} \\ \forall h_1 \in K_1}} \left\langle K \left(g, X, \bigotimes_{h_1} \phi_{h_1, i} \otimes \tilde{\phi}_{h_1, i} \right) \right\rangle \\ & \cdot d(\Omega_{\mathcal{C}})^{|K_1| - |K_2|} \cdot \prod_{h_2 \in K_2} d(X(h_2)). \end{aligned} \quad (5.1)$$

Theorem 5.1. *For a given manifold M , $I_{\mathcal{E}}$ does not depend on the choice of Kirby diagram K , in other words, it is an invariant of smooth, closed, oriented 4-manifolds.*

Proof. We have to prove the following lemmas about $I_{\mathcal{E}}$:

- independence of orientation choices (Lemma 5.4);
- invariance under isotopies of the diagram (Lemma 5.5);
- invariance under handle slides (Lemma 5.7);
- invariance under handle cancellations (Lemma 5.9). ■

Lemma 5.4. *The invariant does not depend on the choice of orientations of the attaching spheres.*

Proof. Let us verify the statement for the attaching spheres of k -handles for all relevant values of k .

(1) A change in orientation of S^0 results in the exchange of the two attaching discs of a 1-handle. In Definition 2.1, we can simply exchange the roles of ι and $\tilde{\iota}$ without changing the diagram.

(2) Changing the orientation of an attaching S^1 of a 2-handle h_2 dualises its labelling object, but also inverts the grade $\deg(h_2)$ of the incident 3-handles (reusing the notation from Definition 5.3). From Remark 4.26, we know that

$$\Omega_{\deg(h_2)}^* \cong \Omega_{\deg(h_2)^{-1}},$$

thus the evaluation is independent of this choice.

(3) Since inversion is an involution on G , we can reindex the summation over 3-handle labellings g and redefine $g(h_3) \mapsto g(h_3)^{-1}$ for the 3-handle h_3 whose attaching sphere was reoriented. ■

Lemma 5.5. *The invariant does not change if an isotopy is applied to a k -handle.*

Proof. Essentially, this is Proposition 4.22, but translated to the labelled planar diagram derived from a Kirby diagram. If the isotopy is regular, the diagram in the pivotal category does not change.

Let us retrace the proof for planar moves occurring from a k -handle isotopy, for any value of k .

(1) Isotopy on a 1-handle (possibly with attachments) may cause point isotopies, which are covered by naturality of the crossed braiding.

(2) The braid, or heptagon, axioms (e.g., [11, (6) and (7)]) and naturality of the crossed braiding ensure invariance under isotopies involving or changing the crossings.

(3) Sliding an attachment of a 3-handle h_3 under a part of the diagram does not change the extracted diagram. Sliding it over a part of the diagram acts on said part

with $g(h_3)^{-1}$, followed by an action with $g(h_3)$. $G \times$ -coherences and their inverses are inserted where the fold arcs of the attaching sphere crosses the remaining diagram. We can cancel these natural isomorphisms and recover the original diagram. (See, e.g., Figure 20, where the two coherences η and η^{-1} can be cancelled.) ■

It is useful to keep the following fact in mind for further proofs and calculations.

Remark 5.6. Recall that G acts on \mathcal{C} via monoidal automorphisms. This implies that on a *closed diagram*, the group action from a 3-handle attachment is trivial, since a closed diagram corresponds to an endomorphism of the monoidal unit \mathcal{I} .

Lemma 5.7. *The invariant does not change if any handle slide is applied to the Kirby diagram.*

Proof. We prove invariance under each j - k -slide.

(1-1) Invariance under this slide is already satisfied for each 3-handle- and 2-handle-labelling individually. It is proven by Corollary C.4 in the appendix, with f the label of the 1-handle attachment being slid, and α the 1-handle to be slid over.

(2-1) Assume, without loss of generality, that the 2-handle S^1 is only slid through halfways. This is again Corollary C.4, where $X = X(h_2)^* \otimes X(h_2)$ and f is the duality coevaluation.

(2-2) This is the graded sliding lemma, 4.31.

(3-1) This does not change the extracted diagram in \mathcal{C} , since attaching a 3-handle to a 1-handle has no effect in the evaluation. (If, in the process of sliding, the 3-handle attachment moves above another part of the diagram, Lemma 5.5 applies.)

(3-2) Up to $G \times$ -coherences, this does not change the extracted diagram.

(3-3) Assume h_3 slides over h'_3 . Up to $G \times$ -coherences, this simply multiplies the labelling $g(h'_3)$ by $g(h_3)$. Therefore, we can reindex the sum over 3-handle-labellings accordingly (since group multiplication is a set isomorphism) and recover the original value.

(3- ∞) Recall Remark 3.13 to visualise this move. Attaching a new 3-handle at infinity that acts on the whole diagram leaves the evaluation invariant (Remark 5.6). Sliding over it is an invariant again, as we just proved. Removing the 3-handle at infinity again does not change the evaluation. ■

In order to prove invariance under handle cancellations, the following lemma is useful.

Lemma 5.8. *Let K_1 and K_2 be Kirby diagrams. $I_{\mathcal{C}}$ is multiplicative under disjoint union of diagrams*

$$I_{\mathcal{C}}(K_1 \sqcup K_2) = I_{\mathcal{C}}(K_1) \cdot I_{\mathcal{C}}(K_2).$$

Proof. Evaluating the disjoint union of two diagrams results in the monoidal product of the evaluations. Since both diagrams are closed, they evaluate to endomorphisms of \mathcal{I} , so the product is simply the multiplication in \mathbb{C} . This shows

$$\langle (K_1 \sqcup K_2)(g, X, \iota) \rangle = \langle K_1(g, X, \iota) \rangle \cdot \langle K_2(g, X, \iota) \rangle,$$

from which the claim follows readily. \blacksquare

Lemma 5.9. *The invariant does not change if cancelling handle pairs are removed from, or added to the Kirby diagram.*

Proof. Perform all necessary handle slides such that the cancelling handle pair is disconnected from the remaining diagram. By Lemma 5.8, we only need to show that the diagram of the cancelling pair evaluates to 1. This is done in the following for the two relevant cancellations.

(1-2) For a simple object $X \in \text{ob } \mathcal{C}$ labelling the 2-handle, the morphism space $\langle X \rangle$ assigned to a 1-handle attaching disc is \mathbb{C} if and only if $X = \mathcal{I}$, and 0 otherwise. In the case it is \mathcal{I} , the basis consists of a single vector, so the sum ranges over a single summand of value 1 (after having noted that the exponent of the normalisation cancels).

(2-3) After consulting Figure 7 and recalling that the closed loop is a diagram for the categorical dimension, it is apparent that the diagram evaluates to $\sum_g d(\Omega_g) = d(\Omega_{\mathcal{C}})$, cancelling the normalisation.

Note that we assume all 0-1-cancellations and 3-4-cancellations to have taken place already. \blacksquare

6. Calculations

As an advantage over state sum models and Hamiltonian formulations (to which the connection will be made in Section 7), it is much easier to calculate explicit values of the invariant. This is mainly because handle decompositions are a very succinct description for smooth manifolds, but also because the graphical calculus of $G \times$ -BSFCs matches Kirby diagrams closely.

This observation was already made in [5] for Kirby diagrams without 3-handles, where the Crane–Yetter invariant was calculated for several 4-manifolds. Due to the following proposition, several results can be recovered for the invariant presented here.

Proposition 6.1. *Let M be a 4-manifold with a handle decomposition that does not contain any 3-handles, and let K be a Kirby diagram for this decomposition. Then*

| Manifold | Invariant value | χ | σ | π_1 |
|---|---|--------|----------|--------------------------------|
| S^4 (incl. exotic candidates) | 1 | 2 | 0 | 1 |
| $S^1 \times S^3$ | $ G \cdot d(\Omega_{\mathcal{C}})$ | 0 | 0 | \mathbb{Z} |
| $S^2 \times S^2, \mathbb{C}\mathbb{P}^2 \# \overline{\mathbb{C}\mathbb{P}^2}$ | $d(\Omega'_e) \cdot d(\Omega_e) \cdot d(\Omega_{\mathcal{C}})^{-2}$ | 4 | 0 | 1 |
| $\mathbb{C}\mathbb{P}^2$ | $\sum_{X \in \mathcal{O}(\mathcal{C}_e)} d(X)^2 \theta_X \cdot d(\Omega_{\mathcal{C}})^{-1}$ | 3 | 1 | 1 |
| $\overline{\mathbb{C}\mathbb{P}^2}$ | $\sum_{X \in \mathcal{O}(\mathcal{C}_e)} d(X)^2 \theta_X^{-1} \cdot d(\Omega_{\mathcal{C}})^{-1}$ | 3 | -1 | 1 |
| $S^1 \times S^1 \times S^2$ (faithful gr.) | $ G ^2 \cdot \mathcal{O}(\mathcal{C}'_e) \cdot d(\Omega_e)$ | 0 | 0 | $\mathbb{Z} \oplus \mathbb{Z}$ |
| $S^1 \times S^3 \# S^1 \times S^3 \# S^2 \times S^2$ | $ G ^2 \cdot d(\Omega'_e) \cdot d(\Omega_e)$ | 0 | 0 | $\mathbb{Z} * \mathbb{Z}$ |

Table 3. Example values for the invariant, in comparison with several classical invariants. χ is the Euler characteristic, σ is the signature, π_1 the fundamental group. $\mathbb{Z} * \mathbb{Z}$ denotes the free product of \mathbb{Z} with itself, i.e., the free group on two generators.

its invariant from Definition 5.3 is equal to its renormalised Crane–Yetter invariant $\widehat{\mathcal{C}\mathbb{Y}}$ [5, Proposition 6.1.1] for the trivial degree, up to a factor involving the Euler characteristic:

$$I_{\mathcal{C}}(M) = \widehat{\mathcal{C}\mathbb{Y}}_{\mathcal{C}_e}(M) \cdot \left(\frac{d(\Omega_{\mathcal{C}})}{d(\Omega_{\mathcal{C}_e})} \right)^{2-\chi(M)}.$$

In particular, if \mathcal{C} is concentrated in the trivial degree, they coincide.

Proof. Without 3-handles, all 2-handles are labelled with Ω_e , and there is no G -action on any part of the diagram. Reminding ourselves that the number of 3-handles $|K_3|$ is zero, furthermore $|K_0| = |K_4| = 1$, and therefore here

$$\begin{aligned} \chi(M) &= |K_0| - |K_1| + |K_2| - |K_3| + |K_4| \\ &= 2 - |K_1| + |K_2|, \end{aligned}$$

we compare to Definition 5.3 and [5, Propositions 4.13 and 6.1.1]. ■

6.1. Example manifolds

The following perspective was pointed out to the author by Ehud Meir: if we fix a particular manifold M^4 and vary \mathcal{C} , the quantity $I_{\mathcal{C}}(M)$ becomes an invariant of $G \times$ -BSFCs. We will show here that many known invariants of fusion categories can be recovered by choosing the correct manifold. Table 3 summarises the results.

6.1.1. $S^1 \times S^3$. Recall the Kirby picture for this manifold from Figure 12. The attaching sphere of the 3-handle, labelled with a group element g , acts trivially in the diagram according to Remark 5.6. There is a single 1-handle and no 2-handle present. Summing over the 1-dimensional morphism space of the monoidal identity

yields

$$I_{\mathcal{C}}(S^1 \times S^3) = \sum_{\substack{g \in G \\ \iota \in (I)}} \left(\iota \right) \left(\tilde{\iota} \right) \cdot d(\Omega_{\mathcal{C}}) = \sum_{g \in G} d(\Omega_{\mathcal{C}}) = |G|d(\Omega_{\mathcal{C}}).$$

6.1.2. $S^2 \times S^2$. Since S^2 admits a standard handle decomposition with a single 0-handle and a single 2-handle, the product handle decomposition of $S^2 \times S^2$ only consists of a 0-handle, two 2-handles and a 4-handle, so Proposition 6.1 is applicable and we can follow the calculation from [5, (6.1.3)] to arrive at

$$I_{\mathcal{C}}(S^2 \times S^2) = \frac{d(\Omega'_{\mathcal{C}_e})d(\Omega_{\mathcal{C}_e})}{d(\Omega_{\mathcal{C}})^2}.$$

For the details of the calculation, assume as Kirby diagram for $S^2 \times S^2$ the Hopf link of two 2-handle attachments [15, Figure 4.30]. Both 2-handles are labelled with $\Omega_{\mathcal{C}_e}$. Then by the killing lemma 4.34, one of them only contributes with the symmetric centre $\Omega'_{\mathcal{C}_e}$. Then the diagram unlinks (since any object in $\Omega'_{\mathcal{C}_e}$ braids trivially with any other object), and we arrive at the desired value after including the normalisation.

6.1.3. Complex projective plane and Gauss sums. The complex projective plane $\mathbb{C}\mathbb{P}^2$ has a handle decomposition with a single 0-handle, 2-handle and 4-handle, so again Proposition 6.1 is applicable and we can reuse the results from [5, Section 3.4]. The 2-handle is attached along the 1-framed unknot (see the equation below), and accordingly the evaluation of this diagram is the trace over the *twist*, whose eigenvalue on a simple, trivially graded object X we will denote by θ_X here:

$$I\left(\bigcirc\right) = \sum_{X \in \mathcal{O}(\mathcal{C}_e)} d(X)^2 \theta_X \cdot d(\Omega_{\mathcal{C}})^{-1}.$$

Flipping the orientation famously results in a non-diffeomorphic manifold, $\overline{\mathbb{C}\mathbb{P}^2}$. Not surprisingly, the Kirby diagram is the mirror diagram of the above, and its invariant is the same with θ replaced by θ^{-1} . The values $\sum_{X \in \mathcal{O}(\mathcal{C}_e)} d(X)^2 \theta_X^{\pm}$ are known as the *Gauss sums* of the ribbon fusion category \mathcal{C}_e .

6.1.4. $S^1 \times S^1 \times S^2$. To study an example where the presence of 3-handles influences the manifold, we turn to $S^1 \times S^1 \times S^2$, borrowing several calculational steps from [5, Section 6.2.1]. For now, we assume that the G -grading on \mathcal{C} is faithful (i.e., $\mathcal{C}_g \not\cong 0$ for any g), which is quite restrictive in the context of $G \times$ -BSFCs; compare Remark 2.4 and note that Corollary 4.36 implies that any faithfully graded $G \times$ -BSFC has natural isomorphisms $X \rightarrow {}^g X$ for every group element G on the transparent subcategory $X \in \mathcal{C}'$. Completing this calculation for non-faithfully graded $G \times$ -BSFCs requires a deeper study of the group action and would probably yield an interesting invariant of the category.

Recall the Kirby diagram from Figure 13 and note that h_{2a} does not have any 3-handle attached to it and will thus be labelled \mathcal{C}_e , while h_{2b} has both 3-handles attached twice with opposite boundary orientation and will be labelled by $\Omega_{[g_1, g_2]}$, where $[g_1, g_2] := g_1 g_2 g_1^{-1} g_2^{-2}$. The number of 1-handles equals the number of 2-handles, cancelling the normalisation, and we can begin to evaluate the diagram

$$I(S^1 \times S^1 \times S^2) = \sum_{\substack{g_1, g_2 \in G \\ X \in \mathcal{O}(\mathcal{C}_e) \\ \alpha, \beta}} d(X) \cdot \begin{array}{c} \begin{array}{ccc} \begin{array}{c} \circlearrowleft \\ \Omega_{[g_1, g_2]} \end{array} & \begin{array}{c} \xrightarrow{g_1 g_2 X} \\ \xrightarrow{g_2 g_1 X} \end{array} & \begin{array}{c} \circlearrowleft \\ \alpha \end{array} \\ \begin{array}{c} \downarrow^{g_2 X} \\ \tilde{\alpha} \end{array} & \begin{array}{c} \xrightarrow{X} \\ \xrightarrow{X} \end{array} & \begin{array}{c} \downarrow^{g_1 X} \\ \tilde{\beta} \end{array} \end{array} \end{array} \quad (6.1)$$

We are using the dual basis convention from Figure 1 when summing over α and β . The contribution from h_{2b} is an encircling where Lemma 4.35 is applicable, so we can restrict the sum over X to \mathcal{C}'_e , since $\mathcal{C}_{[g_1, g_2]} \neq 0$. The morphism spaces of the form $\langle (g_2 X)^* \otimes X \rangle \ni \tilde{\alpha}$ may not be familiar at first, but recall that Corollary 4.36 defines a canonical isomorphism $X \rightarrow g_2 X$, the encircling by Ω_{g_2} . The morphism space $\langle X^* \otimes X \rangle$ is inhabited by the coevaluation $\text{coev}_X: \mathcal{I} \rightarrow X^* \otimes X$, and $\langle (g_2 X)^* \otimes X \rangle$ is spanned by the coevaluation precomposed with the mentioned isomorphism. We thus insert the encircling with Ω_{g_2} graphically where $g_2 X$ enters $\tilde{\alpha}$, and insert the inverse of the isomorphism (which is encircling with Ω_{g_2} and prefactors detailed in Corollary 4.36) at α . The analogous computation can be done for β and $\tilde{\beta}$. These four new encirclings can be slid off the original encircling from h_{2b} , by the sliding lemma (Lemma 4.31). This changes its grade to \mathcal{C}_e , and unlinks the encirclings from the diagram. They can thus be evaluated and result in dimensional factors. Including all prefactors, we have

$$\begin{array}{c} \begin{array}{ccc} \begin{array}{c} \circlearrowleft \\ \Omega_e \end{array} & \begin{array}{c} \xrightarrow{X} \\ \xrightarrow{X} \end{array} & \begin{array}{c} \circlearrowleft \\ \alpha \end{array} \\ \begin{array}{c} \downarrow^X \\ \tilde{\alpha} \end{array} & \begin{array}{c} \xrightarrow{X} \\ \xrightarrow{X} \end{array} & \begin{array}{c} \downarrow^X \\ \tilde{\beta} \end{array} \end{array} \\ = \sum_{\substack{g_1, g_2 \in G \\ X \in \mathcal{O}(\mathcal{C}'_e) \\ \alpha, \beta}} d(X) \cdot \\ \cdot d(\Omega_{g_1}) d(\Omega_{g_2}) d(\Omega_{g_1^{-1}}) d(\Omega_{g_2^{-1}}) \\ \cdot d(X)^{-2} d(\Omega_e)^{-2} d(\Omega_{g_1})^{-1} d(\Omega_{g_1})^{-1}. \end{array}$$

By Lemma 4.27, all factors of the form $d(\Omega_g)$ are equal if they are nonzero. The remaining calculation follows [5, Section 6.2.1]. The ordinary killing lemma (Lemma 4.34) is applied, and evaluations and coevaluations with the correct prefactors inserted for the dual bases:

$$= |G|^2 \cdot |\mathcal{O}(\mathcal{C}'_e)| \cdot d(\Omega_e).$$

Variante: trivial grading on \mathbb{Z}_3 . To see the effect of the group action in a non-faithfully graded $G \times$ -BSFC \mathcal{C} , we choose explicit examples for the category. First, set $G = \mathbb{Z}_2$ and $\mathcal{C}_{\mathbb{Z}_3}^{\mathbb{Z}_2} = \text{Vect}_{\mathbb{Z}_3}$, which denotes \mathbb{Z}_3 -graded finite-dimensional vector spaces, with simple objects $\mathcal{I}, \omega, \omega^*$. Let the grading be concentrated in the trivial degree, and equip the category with the trivial \mathbb{Z}_2 -action, and the trivial braiding. Second, define $\tilde{\mathcal{C}}_{\mathbb{Z}_3}^{\mathbb{Z}_2}$ with the same data as $\mathcal{C}_{\mathbb{Z}_3}^{\mathbb{Z}_2}$, but alter the group action such that the nontrivial element $\sigma \in \mathbb{Z}_2$ acts as ${}^\sigma\omega = \omega^*$, and ${}^\sigma\omega^* = \omega$.

For $\mathcal{C}_{\mathbb{Z}_3}^{\mathbb{Z}_2}$ we jump ahead slightly to Proposition 6.8 and learn:

$$\begin{aligned} I_{\mathcal{C}_{\mathbb{Z}_3}^{\mathbb{Z}_2}}(S^1 \times S^1 \times S^2) &= |\{\phi: \mathbb{Z} \oplus \mathbb{Z} \rightarrow \mathbb{Z}_2\}| \cdot \widehat{\text{CY}}_{\text{Vect}_{\mathbb{Z}_3}}(S^1 \times S^1 \times S^2) \\ &= 4 \cdot 9 = 36. \end{aligned}$$

For $\tilde{\mathcal{C}}_{\mathbb{Z}_3}^{\mathbb{Z}_2}$, we have to follow the calculation from (6.1). Since \mathbb{Z}_3 is abelian, we get $\Omega_{[g_1, g_2]} = \Omega_e$ throughout. The encircling unlinks and contributes as a global factor of $d(\Omega_e) = 3$. We do not have natural isomorphisms $X \cong {}^g X$ and need to perform the sums over the morphisms explicitly. Luckily, the morphism spaces $\langle {}^g X^* \otimes X \rangle$ are only \mathbb{C} for $g = e$ or $X = \mathcal{I}$, and 0 otherwise. The diagram evaluates to 1 if both morphism spaces are \mathbb{C} , which leaves us to merely count the admissible labels (X, g_1, g_2) :

$$I_{\tilde{\mathcal{C}}_{\mathbb{Z}_3}^{\mathbb{Z}_2}} = |\{(\mathcal{I}, g_1, g_2) | g_1, g_2 \in \mathbb{Z}_2\} \sqcup \{(\omega, e, e), (\omega^*, e, e)\}| \cdot d(\Omega_e) = 6 \cdot 3 = 18.$$

6.2. Connected sum, smooth structures, and simply-connectedness

It is a natural question to ask how strong the invariant is, and in particular whether it can detect smooth structures. In this subsection, we demonstrate that generically, this is not the case.

Proposition 6.2. *The invariant is multiplicative under connected sum. Explicitly, let M_1 and M_2 be two manifolds, and $M_1 \# M_2$ their connected sum. Then,*

$$I_{\mathcal{C}}(M_1 \# M_2) = I_{\mathcal{C}}(M_1) \cdot I_{\mathcal{C}}(M_2).$$

Proof. Given two Kirby diagrams for M_1 and M_2 , respectively, their disjoint union is a diagram for $M_1 \# M_2$. Then the statement follows from Lemma 5.8. ■

This proposition has far-reaching consequences for generic $G \times$ -BSFCs.

Corollary 6.3. *Assume furthermore that the Gauss sum of \mathcal{C}'_e is invertible. Then $I_{\mathcal{E}}$ does not detect exotic smooth structures.*

Proof. By [13, Proposition 6.11], the Gauss sums of \mathcal{C}_e are invertible, and thus also $I_{\mathcal{E}}(\mathbb{C}\mathbb{P}^2)$ and $I_{\mathcal{E}}(\overline{\mathbb{C}\mathbb{P}^2})$, by the results of Section 6.1.3. We can apply the remark from [20, Section 1.4] to infer that $I(M)$ for any M depends only on the signature, the Euler characteristic, the fundamental group and the fundamental class of M . ■

Corollary 6.4. *Assume again that the Gauss sum of \mathcal{C}'_e is invertible. For a simply-connected manifold M with Euler characteristic $\chi(M)$ and signature $\sigma(M)$, the invariant readily computes as*

$$I_{\mathcal{E}}(M) = I_{\mathcal{E}}(\mathbb{C}\mathbb{P}^2)^{\frac{\chi(M)+\sigma(M)}{2}-1} \cdot I_{\mathcal{E}}(\overline{\mathbb{C}\mathbb{P}^2})^{\frac{\chi(M)-\sigma(M)}{2}-1}. \quad (6.2)$$

Proof. See, e.g., [5, Lemma 3.12]. ■

Example 6.5. Under the same assumptions as before, we have $I_{\mathcal{E}}(\mathbb{C}\mathbb{P}^2 \# \overline{\mathbb{C}\mathbb{P}^2}) = I_{\mathcal{E}}(S^2 \times S^2)$.

Remark 6.6. The invariant in its generic form does not seem to be helpful in the search for exotic smooth 4-manifolds, nor does it depend directly on the 3-type of the manifold (as was hoped, e.g., in [32, Section 5]).

In line with what we noted at the beginning of Section 6.1, the more fruitful viewpoint is to study the resulting invariants for $G \times$ -BSFCs when fixing a particular manifold. By Table 3, we could recover the global dimensions of \mathcal{C} , \mathcal{C}_e and \mathcal{C}'_e , as well as the rank of \mathcal{C}'_e , and the Gauss sums. By the next subsection we will be able to recover most information about the group G . It remains an interesting open question whether there is a finite set of manifolds such that the corresponding set of invariants is complete on $G \times$ -BSFCs, i.e., can distinguish any two $G \times$ -BSFCs up to equivalence.

6.3. Untwisted Dijkgraaf–Witten theory

Definition 6.7. Let G be a finite group. The *untwisted Dijkgraaf–Witten invariant* $DW_G(M)$ is defined as the number of G -connections on the manifold M :

$$DW_G(M) := |\{\phi: \pi_1(M) \rightarrow G\}|.$$

This normalisation is not the most common in the literature, but it is simpler for our purposes.

Proposition 6.8 (Compare [11, Proposition 4.5]). *Let \mathcal{C} be concentrated in the trivial degree, i.e., $g \neq e \implies \mathcal{C}_g \simeq 0$ (or equivalently $\mathcal{C} \simeq \mathcal{C}_e$ as fusion categories), and the G -action be trivial. Then the invariant is a product of the Dijkgraaf–Witten invariant and the renormalised Crane–Yetter invariant [5, Proposition 6.1.1]:*

$$I_{\mathcal{C}}(M) = DW_G(M) \cdot \widehat{CY}_{\mathcal{C}_e}(M).$$

Proof. Any G -labelling on 3-handles for which any 2-handle is labelled with Ω_g such that $g \neq e$ does not contribute to the invariant since $\Omega_g = 0$ in that case. Recall the presentation of the fundamental group from Section 3.4.3. The 3-handles are generators of $\pi_1(M)$, while the 2-handles give relations. A homomorphism from $\pi_1(M)$ to G is presented by an assignment of a G -element for every 3-handle, such that the relations of $\pi_1(M)$ are satisfied. The former is given by a G -labelling, the latter is implemented by the fact that only those G -labellings contribute to the invariant where all 2-handles are labelled by Ω_e . Thus, the contributing G -labellings indeed run over all possible homomorphisms $\phi: \pi_1(M) \rightarrow G$. As in the proof of Proposition 6.1, it is easy to see that each such contribution is just $\widehat{CY}_{\mathcal{C}_e}(M)$. ■

7. Recovering the state sum model

The main motivation for this work was the second open question in [11, Section 7]. There, a state sum model is defined for $G \times$ -BSFCs in terms of triangulations, but it is noted that the model is impractical to compute. We will show now that the invariant defined here is equal to the state sum up to a classical factor, giving an economical way to calculate the state sum.

The state sum takes a $G \times$ -BSFC \mathcal{C} and a manifold M with compatible triangulation \mathcal{T} , labels 1-simplices with group elements, 2-simplices with simple objects and 3-simplices with morphisms. For every 4-simplex, a local partition function is defined.

There is a well-known procedure called *chain mail* to turn an invariant defined on Kirby diagrams into a state sum model. A smooth triangulation gives rise to a handle decomposition, and we have to pull back the definition along this procedure. Chain mail has been described already in [26, Section 4.3]. As we go through the steps, we refer to [5, Section 5] for details, which matches our conventions and notation to a large degree.

Definition 7.1 (Well known). Every triangulation \mathcal{T} of a smooth manifold M gives rise to a handle decomposition, where the k -simplices are thickened to $(n - k)$ -handles. This is called the *dual handle decomposition*. The set of k -simplices is denoted as \mathcal{T}_k .

Theorem 7.1. *Up to a factor involving the Euler characteristic, the invariant from Definition 5.3 is equal to the state sum $Z_{\mathcal{E}}$ from [11, (23)]. Explicitly, let M be a connected manifold and \mathcal{T} an arbitrary triangulation, then*

$$I_{\mathcal{E}}(M) = Z_{\mathcal{E}}(M; \mathcal{T}) \cdot d(\Omega_{\mathcal{E}})^{1-\chi(M)} \cdot |G|.$$

Proof. We begin with the dual handle decomposition of \mathcal{T} . It does not have a single 0-handle and 4-handle, thus we cannot directly calculate I of this decomposition. The strategy is to slightly modify the manifold until we can calculate the invariant, and recover the original value from there. Each 0-handle can be regarded as a drawing canvas which contains the Kirby diagram of the corresponding 4-simplex σ . Proposition 7.2 will show that the evaluation of such a diagram (denoted here as $\langle \sigma \rangle$) is just the quantity $Z_F^{\pm}(\sigma)$ from [11, Section 3.1].

In order to join all drawing canvases by cancelling the excessive 0-handles, we attach $|\mathcal{T}_4| - 1$ 1-handles (recall that the handle decomposition had $|\mathcal{T}_4|$ 0-handles to begin with). Since the resulting manifold is not closed, we attach $|\mathcal{T}_4| - 1$ 3-handles to cancel the boundary. The resulting manifold is $N := M \#^{|\mathcal{T}_4|-1} S^1 \times S^3$, for which we know from Proposition 6.2

$$I_{\mathcal{E}}(M) = I_{\mathcal{E}}(N) \cdot |G|^{1-|\mathcal{T}_4|} d(\Omega_{\mathcal{E}})^{1-|\mathcal{T}_4|}.$$

We know that still $|\mathcal{T}_0| - 1$ 3-handles will be cancelled by excessive 4-handles, then we can calculate $I_{\mathcal{E}}(N)$ diagrammatically. Since the diagram for N is a disjoint union of diagrams for each 4-simplex $\sigma \in \mathcal{T}_4$ (with $|\mathcal{T}_4| - 1$ 3-handles added and $|\mathcal{T}_0| - 1$ removed, each incurring a factor of $|G|$), it factors as a product like in the proof of Lemma 5.8, and we can infer from (5.1)

$$I_{\mathcal{E}}(N) = \sum_{g, X} d(\Omega_{\mathcal{E}})^{|\mathcal{T}_3|-|\mathcal{T}_2|} |G|^{|\mathcal{T}_4|-|\mathcal{T}_0|} \prod_{h_2} d(X_{h_2}) \prod_{\sigma \in \mathcal{T}_4} \langle \sigma \rangle.$$

We compare with [11, (23)], which uses the notation $D^2 = d(\Omega_{\mathcal{E}})$ and labels objects with f , and defines (in our notation)

$$Z_{\mathcal{E}}(M; \mathcal{T}) = \sum_{g, X} d(\Omega_{\mathcal{E}})^{|\mathcal{T}_0|-|\mathcal{T}_1|} |G|^{-|\mathcal{T}_0|} \prod_{h_2} d(X_{h_2}) \prod_{\sigma \in \mathcal{T}_4} \langle \sigma \rangle.$$

We combine all three equations and recall that the Euler characteristic is $\chi(M) = |\mathcal{T}_0| - |\mathcal{T}_1| + |\mathcal{T}_2| - |\mathcal{T}_3| + |\mathcal{T}_4|$ to verify the claim. \blacksquare

Proposition 7.2. *The Kirby diagram of a 4-simplex σ evaluates to the quantity $Z_F^{\pm}(\sigma)$ from [11, Section 3.1].*




Figure 22. The Kirby diagram of a 4-simplex, showing only the nontrivially acting 3-handle.

Proof. The only novelty over [26, Section 4.3] and [5, Section 5] is the appearance of 3-handles in the Kirby diagram of a 4-simplex. Each 3-handle is given by a 1-simplex, or edge, in the triangulation, and is thus specified by its two end vertices. For any third vertex (of which there exist three), a 2-simplex, or triangle, exists such that the edge is part of the boundary of the triangle, and thus the 3-handle attaches to the corresponding 2-handle. There are five 3-simplices (each opposite a vertex) which are thickened to five 1-handles. These connect the 4-simplex to any neighbouring 4-simplex, and thus only one of the attaching D^3 is visible in the diagram. Since a 3-simplex has four boundary triangles, four 2-handles attach to each 1-handle. The resulting diagram is shown in Figure 22, but not all 3-handles are shown, for clarity. The diagram is not unique since any 3-dimensional isotopy can be applied to it, but it is favourable to minimise the number of crossings to keep calculations simple. Similarly, it is desirable to minimise the number of 3-handles covering a 1-handle attaching disc. The minimum number is one each, and it is achieved in the diagram, following the convention of [11, Figure 13 (left)]. ■

7.1. From state sum models to TQFTs

There is a well-known recipe to define a TQFT $\mathcal{Z}_{\mathcal{E}}$ from a topological state sum model [31] such that $\mathcal{Z}_{\mathcal{E}}(M^4) = Z_{\mathcal{E}}(M^4)$ for a 4-manifold, where M^4 is regarded as a cobordism from the empty manifold to itself. The dimensions of the state spaces

$\mathcal{Z}_{\mathcal{C}}(N^3)$ assigned to boundary 3-manifolds are of interest and can be calculated as $\dim \mathcal{Z}_{\mathcal{C}}(N^3) = \text{tr } 1_{\mathcal{Z}_{\mathcal{C}}(N^3)} = \mathcal{Z}_{\mathcal{C}}(S^1 \times N^3) = \mathcal{Z}_{\mathcal{C}}(S^1 \times N^3)$. Using Theorem 7.1, this calculation is now much easier than directly from the state sum.

Corollary 7.3. *The dimension of the state space assigned to a 3-manifold N is*

$$\dim(\mathcal{Z}_{\mathcal{C}}(N)) = I_{\mathcal{C}}(S^1 \times N) \cdot d(\Omega_{\mathcal{C}})^{-1} \cdot |G|^{-1}.$$

Proof. Note that the Euler characteristic is multiplicative and $\chi(S^1 \times M) = \chi(S^1) = 0$. We continue the calculation of $\mathcal{Z}_{\mathcal{C}}(S^1 \times M)$ to arrive at the result. ■

Examples 7.4. (1) One readily verifies that $\dim(\mathcal{Z}_{\mathcal{C}}(S^3)) = 1$. It is well known that such a TQFT is, up to a factor of $\mathcal{Z}_{\mathcal{C}}(S^4)^{-1}$, multiplicative under connected sum, in agreement with Proposition 6.2.

(2) For a faithfully graded $G \times$ -BSFC, we calculate

$$\dim(\mathcal{Z}_{\mathcal{C}}(S^1 \times S^2)) = |G| \cdot |\mathcal{O}(\mathcal{C}'_e)| \cdot d(\Omega_e) \cdot d(\Omega_{\mathcal{C}})^{-1} = |\mathcal{O}(\mathcal{C}'_e)|,$$

showing that the theory is non-invertible when \mathcal{C}_e is not modular.

(3) For the two examples from \mathbb{Z}_2 acting on $\text{Vect}_{\mathbb{Z}_3}$, we get

$$\dim(\mathcal{Z}_{\mathcal{C}_{\mathbb{Z}_2}^{\mathbb{Z}_3}}(S^1 \times S^2)) = 6 \quad \text{and} \quad \dim(\mathcal{Z}_{\mathcal{C}_{\mathbb{Z}_3}^{\mathbb{Z}_2}}(S^1 \times S^2)) = 3.$$

7.2. Discussion: defining a TQFT directly from handle decompositions

Manifolds with nonempty boundary are described by handle decompositions as well [15, Section 5.5]. In fact, handle attachments can be seen as the fundamental generators of bordisms, and TQFTs can be defined naturally in terms of them [19].

An arbitrary Kirby diagram (one that does not necessarily correspond to a closed manifold) specifies both a boundary 3-manifold N and a bordism $M: \emptyset \rightarrow N$. The boundary N is constructed via *surgery*, as per Remark 3.3. It is in principle possible to define a TQFT directly from handle decompositions of bordisms, but doing so rigorously is beyond the scope of this article and is left as future work.

In Atiyah's axiomatisation of TQFTs [2], a vector space $\mathcal{Z}(N)$ is assigned to every boundary manifold N , and a vector $\mathcal{Z}(M) \in \mathcal{Z}(\partial M)$ is assigned to every top-dimensional manifold. We informally propose these two constructions for a $G \times$ -BSFC \mathcal{C} . (It should also be possible to generalise the approach to functorial TQFTs, and repeat the constructions for bordisms $M: N_1 \rightarrow N_2$ with a nonempty domain N_1 by using relative Kirby calculus.)

In analogy to the Turaev–Viro model [22], one would define a *string net space*, or “skein space,” for $\mathcal{Z}(N)$. Essentially, it would be defined as the vector space over \mathcal{C} -labelled diagrams (as in Definition 4.15) *embedded in N* , modulo local relations

in \mathcal{C} (the evaluation from Definition 4.16). For S^3 , this space is then tautologically 1-dimensional since the labelled diagrams are defined to be evaluated on it. But for a more complicated manifold N , the dimension of its state space should be higher, and should in particular coincide with the values $I_{\mathcal{C}}(S^1 \times N)$ derived in the previous subsection.

For a given Kirby diagram K , the manifold $S^3(K)$ is defined as surgery on S^3 along the spheres in K . The string net space $\mathcal{Z}(S^3(K))$ has then an easier description as \mathcal{C} -labelled diagrams embedded in $S^3 \setminus K$, modulo local relations and Kirby moves. (The latter are a complete set of moves that relate any two surgery diagrams of diffeomorphic 3-manifolds.)

The vacuum state assigned to $S^3(K)$ (i.e., the vector corresponding to $M: \emptyset \rightarrow N$, where K is a Kirby diagram for M) should then simply be the empty diagram. This could confuse at first since different M potentially give rise to different vacua, but note that the vacuum state is only the empty diagram in the particular surgery diagram K , and would usually not be mapped to the vacuum state of $S^3(K')$ under a diffeomorphism $S^3(K) \cong S^3(K')$, for a different diagram K' corresponding to a non-diffeomorphic bordism $M': \emptyset \rightarrow N$.

8. Spherical fusion 2-categories and related work

This section is kept in an informal discussion style. One reason to adopt the graphical calculus presented in Section 4 was to stay close to the language of Kirby diagrams with 3-handles, but another reason was to imitate the graphical calculus of Gray categories (semistrict 3-categories) with duals [7]. The main reason not to translate Kirby diagrams into Gray category diagrams is to take advantage of the graphical calculus of pivotal categories, which is a considerable shortcut.

$G \times$ -BSFCs can be seen as monoidal 2-categories [11, 12], which in turn can be seen as one-object 3-categories, this approach is thus natural. The translation of $G \times$ -BSFCs into monoidal 2-categories is summarised in Table 4.

Since [23], there is a search for a good categorification of spherical fusion categories to the world of monoidal 2-categories. The definition proposed there was soon found to be too restrictive, but it was long an open problem to find a more general and still well-behaved definition. It is shown in [11] explicitly that $G \times$ -BSFCs already yield more general monoidal 2-categories, but still a good definition of spherical fusion 2-category is expected to be much more general.

[12] offers a detailed and promising definition, relates it to [11], and defines a state sum model. It is to be expected that the invariant based on Kirby diagrams with 3-handles presented here can be generalised to that notion of spherical fusion 2-categories, and that it will coincide with that state sum. Some elements of this art-

| Datum in a $G \times$ -BSFCs | Notation | Datum in the monoidal 2-category |
|------------------------------|---|----------------------------------|
| Group element | $g \in G$ | Object |
| Grade | \mathcal{C}_g | Endocategory |
| Object | $A \in \mathcal{C}$ | 1-morphism |
| Morphism | $f: A \rightarrow B$ | 2-morphism |
| Group multiplication | $g_1 g_2$ | Monoidal product |
| Group inverse | g^{-1} | Duality |
| Group action | ${}^s A$ | 1-morphism composition |
| Monoidal product | $A \otimes B$ | 1-morphism composition |
| Crossed braiding | $c_{X,Y}: X \otimes Y \rightarrow {}^s Y \otimes X$ | Interchanger |

Table 4. The translation of $G \times$ -BSFCs to monoidal 2-categories. Summarised from [11, Section 6] and [12, Constructions 2.1.23 and 2.3.6].

icle that played no role in the invariant, such as the fold lines of sheets, are then less trivial to handle. The details of this correspondence are left for future work, but one can speculate that the handle picture will once again allow for much more efficient computations.

In the following discussion, we want to informally discuss such a generalisation, and digress on a good “higher spherical axiom.”

4-cocycles and pentagonators. The Dijkgraaf–Witten model is usually defined for the datum of a finite group G and a 4-cocycle $\omega \in C^4(G, U(1))$ (with cohomologous cocycles giving rise to equivalent theories). If $[\omega] \neq 0$, the theory is called *twisted*, but in Section 6.3, only the untwisted Dijkgraaf–Witten model occurred. In [11, Section 4.4], the state sum model is generalised to include a 4-cocycle, and the twisted Dijkgraaf–Witten model is recovered.

The reason seems to be that the monoidal 2-category defined from a $G \times$ -BSFC is very strict. In a general monoidal bicategory (see, e.g., [27, Appendix B.4]), there exist associators for the monoidal product, and even these do not satisfy the pentagon axiom on the nose, but rather up to an invertible modification, the *pentagonator*. In a monoidal 2-category from a $G \times$ -BSFC, the monoidal product is associative though, and the pentagonator is trivial. This seems quite restrictive, and we would want to allow for a weaker structure.

Still with an associative monoidal product, the pentagonator can contain nontrivial data. It consists of an invertible automorphism of the identity 1-morphism for the tensor product of every four objects, satisfying the “associahedron equation.” For the case of monoidal 2-categories from $G \times$ -BSFCs, one can verify that pentagonators in fact correspond to 4-cocycles of G . It would be interesting to study to what data all

further coherences of monoidal 2-categories correspond, and to define the invariant, or the state sum model, in terms of it.

Note though that in [12, Definition 3.3.8], no pentagonator occurs as structure of the monoidal category, but it is strictified, as in Gray categories.

The 3-spherical axiom. The 3-dimensional Turaev–Viro–Barrett–Westbury model ([8,31]) is most generally defined for a *spherical* fusion category. The spherical axiom demands that the left and right traces defined by the pivotal structure are equal. Given a diagram in \mathbb{R}^2 of a morphism in a spherical fusion category, this axiom is equivalent to embedding the diagram in $S^2 \cong \mathbb{R}^2 \cup \infty$ and allowing any line to be isotoped past ∞ . Unsurprisingly, this is the lower-dimensional analogon of the 3- ∞ -move from Section 3.3.2: 3-manifolds have handle decompositions, and when depicting them in the plane, there is a “2- ∞ -move” between different diagrams of the same decomposition. The graphical calculus in the category must correspondingly validate this move. (One might argue that the move should thus be called the “2-spherical” move, since it takes place on a 2-sphere.)

It is only natural to require a “3-spherical axiom” in any suitable definition of spherical fusion 2-categories which corresponds to the 3- ∞ -move. In those 2-categories which come from $G \times$ -BSFCs, this move is trivially true [12, Example 2.3.6], since group elements act trivially on the tensor unit of the category.

Spherical fusion 2-categories satisfy precisely this move [12, Definition 2.3.2], justifying its name and its relevance. In this greater generality, it becomes a necessary ingredient to define dimensions of objects of 2-categories. In the graphical calculus, a trivially attached 3-handle (Figure 3) would evaluate to the dimension of its label. In $G \times$ -BSFCs, these dimensions are always 1, and thus irrelevant, but in general, 3-handles would play a bigger computational role.

Defect theories and orbifold data. It seems fruitful to compare to orbifold theories of defect theories. The Turaev–Viro–Barrett–Westbury model is an orbifold TQFT of the trivial 3d defect theory [9], and the parametrising spherical fusion categories are algebras in a tricategory. Gray categories are naturally objects in a (yet to be rigorously defined) tetracategory, and one would hope that the Gray category derived from a $G \times$ -BSFC as well as the orbifold data of the trivial 4d defect theory carry the appropriate algebra structure satisfying the correct higher sphericity axiom.

A. Manifolds with corners

A closed n -manifold is defined to be locally homeomorphic to \mathbb{R}^n , whereas a manifold with boundary is locally homeomorphic to (a neighbourhood in) $\mathbb{R}^{n-1} \times \mathbb{R}_+$, with

$\mathbb{R}_+ := [0, \infty)$. Closed manifolds have products, but defining a product for manifolds with boundary is not straightforward since corners arise.

A (smooth) n -manifold with k -corners is locally homeomorphic to $\mathbb{R}^{n-k} \times \mathbb{R}_+^k$. A closed manifold thus has 0-corners, a manifold with boundary has 1-corners, and the product of two manifolds with boundaries is naturally a manifold with 2-corners.

Since as topological spaces $\mathbb{R}^{n-1} \times \mathbb{R}_+ \cong \mathbb{R}^{n-k} \times \mathbb{R}_+^k$, the notion of corners only becomes relevant in the context of smooth manifolds. The boundary of a smooth n -manifold with k -corners is subtle to define. The naive definition, in which a point is on the boundary whenever any of the last k coordinates is 0, does not yield a smooth manifold, only a topological manifold. In a sensible definition like [18, Definition 2.6], the boundary is a smooth $(n-1)$ -manifold with $(k-1)$ -corners, but it is in general not a submanifold. Rather, j disjoint copies are made for every j -corner point, and the boundary is immersed in the n -manifold. For example, the archetypal corner manifold \mathbb{R}_+^2 has as boundary $\mathbb{R}_+ \times \{0\} \sqcup \{0\} \times \mathbb{R}_+$, where the origin $(0, 0)$ appears twice.

In this light, handles and their diverse regions are much easier to understand conceptually. The k -disc D^k is a manifold with boundary for $k \geq 2$, and thus an n -dimensional k -handle $h_k = D^k \times D^{n-k}$ is generally an n -manifold with 2-corners. Its boundary is the disjoint union $S^{k-1} \times D^{n-k} \sqcup D^k \times S^{n-k-1}$, and we have called the former “attaching region” and the latter “remaining region.” It is an $(n-1)$ -manifold with 1-corners, or simply with boundary.

As a topological manifold, all n -dimensional k -handles are homeomorphic to D^n , reflecting the fact that we can glue the attaching region and the remaining region together along their common boundary $S^{k-1} \times S^{n-k-1}$ to arrive at the sphere S^{n-1} .

Glueing two handles along their boundary regions results in a manifold with corners, but usually these corners are of index 2. There is a canonical way to “smoothen” these corners [15, Remark 1.3.3], so a handle body again becomes a manifold with boundary after the handle attachment.

As a simple example, we define n -gons:

Definition A.1. The n -gon G_n is the unique oriented, simply connected 2-manifold with corners whose boundary ∂G_n consists of n copies of the interval. In particular,

- $G_0 = S^2$ is the 2-sphere;
- G_1 is the “teardrop” manifold (see, e.g., [27, Section 3.1, Figure 1]);
- G_3 is the 2-simplex.

For $n > 0$, the underlying topological manifold is the 2-disc D^2 , but its smooth corner structure is different.

B. Details on Kirby calculus and diagrams

B.1. Kirby convention

Proof of Lemma 3.12. We have to show that any handle decomposition with a single 0-handle and 4-handle can be isotoped such that it is *regular* (Definition 3.7) and satisfies the *single picture conventions* (Definition 3.9).

Regularity can always be achieved by isotoping the attaching regions, see, e.g., [15, Proposition 4.2.7].

Let us show the two conditions for the single picture conventions.

(2-1, 3-1) This has been illustrated in Figure 6. Inside the remaining region $\partial_r h_1 = [-1, +1] \times S^2$ of a 1-handle, apply an isotopy such that all attachments intersect $\{0\} \times S^2$ transversely in a finite subset (for a 2-handle attachment) or a 1-dimensional manifold (for a 3-handle attachment). By transversality, this can be extended to a small neighbourhood $(-\varepsilon, +\varepsilon) \times S^2$, and then pushed out of $\partial_r h_1$ with an isotopy into the main drawing canvas.

(3-2) Analogously to the situation for 1-handles, focus on $\{0\} \times S^1 \subset D^2 \times S^1 = \partial_r h_2$ inside the remaining region of a 2-handle, and isotope any 3-handle attachment such that it intersects this circle transversely in a finite set of points. Extend to the thickening $D_\varepsilon^2 \times S^1$ (where D_ε^2 is a disc around 0 with radius ε) and isotope out of the remaining region by enlarging this disc, possibly pushing all folds outside into the drawing canvas. ■

B.2. Blackboard framing

An embedded thickened ribbon is given, up to isotopy, by a framed ribbon, which is a ribbon with a nonvanishing normal vector field. Immersing an oriented ribbon into an oriented surface defines a framing by arbitrarily choosing a vector field into the left-hand side of the surface, as viewed from the ribbon direction.

A regular diagram as in Definition 4.7 defines a blackboard framing on its ribbons and fold arcs. Similarly, a Kirby diagram defines a blackboard framing for its 2-handle attachments. For any given framing, it is always possible to match it with the blackboard framing by repeatedly applying the first Reidemeister move.

The notion of blackboard framing is standard for Kirby diagrams without 3-handles, but has not been described yet for diagrams with 3-handles. The projection of the diagram onto the plane locally defines a right-hand side and a left-hand side for every ribbon. To assume the blackboard framing for at 3-handle attachment means to require that it may not change the side, so the 3-handle attaching sphere must stay parallel to the plane close to where it attaches to a 2-handle.

As with any convention, it must be shown that later constructions do not depend on the choices made. Here, we have to show that the invariant does not depend on whether

a particular 3-handle is attached on the right-hand side or the left-hand side. Assume that 3-handles labelled g_1, g_2, \dots, g_K are attached to a 2-handle on the right-hand side (from top to bottom) and g_{K+1}, \dots, g_N, g on the left-hand side (from bottom to top), which implies that the 2-handle is labelled with $\Omega_{g'g}$, where $g' := g_1 g_2 \cdots g_N$. If the 3-handle labelled g is isotoped around to the right side, the grade changes from $g'g$ to gg' , and the 3-handle attaching sphere now covers the 2-handle circle. Since ${}^g\Omega_{g'g} \cong \Omega_{gg'}$, the diagram still evaluates to the same quantity.

B.3. Related graphical calculi

A standard reference for graphical calculus is [29, Sections 3 and 4]. Like the one presented here, it also resembles the calculus of braided categories, but instead of using 2-dimensional sheets, it implements the G -crossed structure by means of a group homomorphism from the diagram complement to G . We chose a slightly different route in this article and defined a more elaborate diagram language which matches Kirby calculus with 3-handles more closely, but is also similar to the graphical calculus of semistrict 3-categories with duals [7], as mentioned in Section 8. Turaev's calculus can be recovered soundly from the one presented here in Section 4, as will be shown now.

First, we will paraphrase the relevant definitions from [29, Sections 3 and 4] in our conventions.

Definition B.1. Let G be a group, and \mathcal{C} a $G \times$ -BSFC. A \mathcal{C} -coloured G -link consists of

- a chosen base point $z \in S^3$,
- a framed link $L \in S^3$,
- a group homomorphism $g: \pi_1(S^3 \setminus L) \rightarrow G$,
- for every homotopy class of paths γ from z to some point in L , an object X_γ in \mathcal{C} .

The objects X_γ are subject to certain coherence laws detailed in [29, Section 3.1]. For example, the action of a loop $\beta: S^1 \rightarrow S^3 \setminus L$ on a path γ corresponds to the group action of $g([\beta])$ on the labelling object X_γ .

Whenever a diagram (Definition 4.1) is in particular a link, a \mathcal{C} -colouring can be recovered from a labelling (Definition 4.15).

Definition B.2. Let \mathcal{D} be a labelled diagram with no points. Construct the associated \mathcal{C} -coloured G -link as follows.

- The base point is ∞ .
- The link is given by the ribbons \mathcal{D}_1 , which necessarily are all circles.

- The group homomorphism is defined on a loop around a circle as the product of group elements of the incident sheets on the circle, just as in the typing relations of 4.15.
- For each circle r , choose one path from ∞ to it, running below the diagram in the projection onto the plane. Its homotopy class is assigned the labelling object $X(r)$. Other homotopy classes of paths to the same circle are assigned objects according to the conditions in [29, Section 3.1].

It is a tedious exercise to see that the group homomorphism and the path colorings are indeed well defined.

The calculus of \mathcal{C} -coloured links is generalised in [29, Section 4] to graphs with coupons labelled with morphisms in the category. The definition can then be repeated with general labelled diagrams, i.e., with points.

Lemma B.3. *The evaluation of labelled diagrams (Definition 4.16) extends to a functor of $G \times$ -BSFCs.*

Proof. This is a routine generalisation where the definition of diagrams has to be extended to include ribbons with open ends. Functoriality then follows directly from the graphical calculus of pivotal fusion categories. ■

Corollary B.4. *Given a labelled diagram, its evaluation (Definition 4.16) coincides with the action of the unique functor defined in [29, Theorems 3.6 and 4.5] on the associated \mathcal{C} -coloured graph.*

Proof. The translation of the ribbons to (crossed) braidings is well known, see [30, Chapter I] and [29, Theorem 4.5]. It remains to check whether the group action evaluates the same. The group action of the calculus in this work differs by the one from [29] just by the extra flexibility arising from folds of the sheets. This gives rise to isomorphic, but different group actions, and thus to different placements of the coherence isomorphisms. An example is Figure 20. But from the coherence theorem of $G \times$ -BSFCs [11, 24] we know that the whole morphism, and thus the evaluation is invariant with respect to these choices. ■


C. 1-handle slide lemmas in fusion categories

To the knowledge of the author, the following lemmas are not mentioned explicitly in the literature.

Definition C.1. Let \mathcal{C} be a pivotal fusion category and $A \in \mathcal{C}$ an object. Assuming the summation conventions from Figure 1, the *projection onto the monoidal unit*

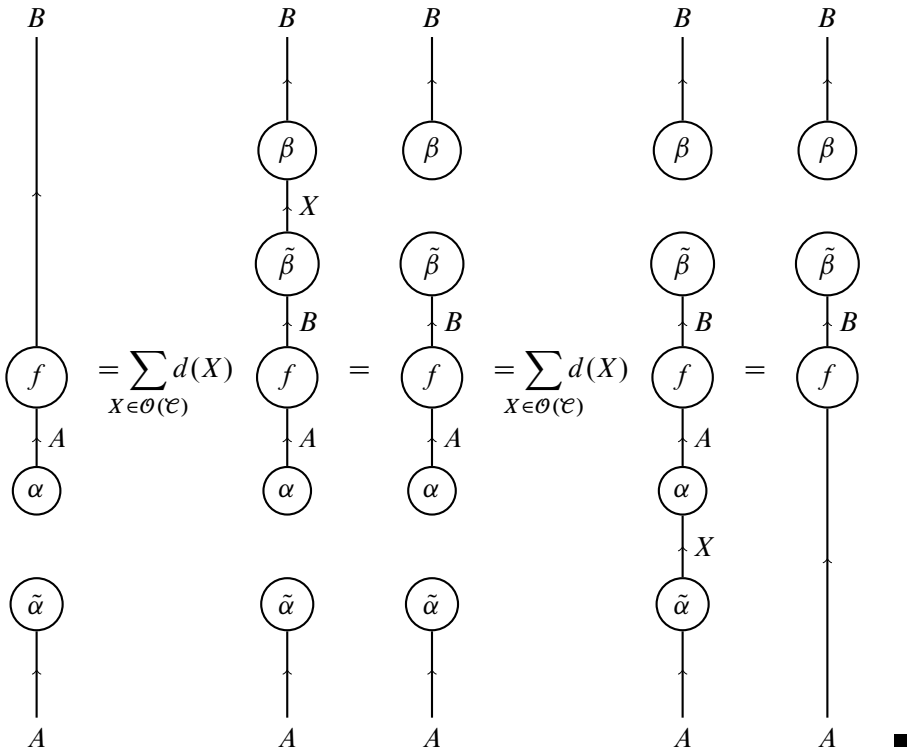
$$\pi_A^I: A \rightarrow A$$

is defined as the following morphism:




Lemma C.2. *The projection onto the monoidal unit is a natural transformation, i.e., π_A^I is natural in A .*

Proof. Explicitly, let $f : A \rightarrow B$ be a morphism in \mathcal{C} . Then the following is valid:



Remark C.3. For $G \times$ -BSFCs, the identity holds even if there is a group action on $\tilde{\alpha}$ or α . In the middle diagram, the disconnected part $\alpha \circ f \circ \tilde{\beta}$ forms a closed diagram, and by Remark 5.6 it can be moved freely from into or out of a group action. ■

Corollary C.4. *The 1-1-handle slide is valid in the graphical calculus of pivotal fusion categories. Explicitly, let $f: \mathcal{I} \rightarrow A$. Then the following is true:*



Proof. Apply the previous lemma to $f \otimes 1_B$. ■

Acknowledgements. For helpful discussions and correspondence, thank you, Daniel Scherl, Gregor Schaumann, Nils Carqueville, Flavio Montiel Montoya, Sascha Biberhofer, Dominic Williamson, Alex Bullivant, Bruce Bartlett, Andras Juhász, Dominik Wrzadzidlo, Chris Douglas, David Reutter, Shawn Cui, Noah Snyder, Marcel Bischof, Arun Debray, and the anonymous reviewer.

Funding. The author wishes to thank the Peters-Beer-Stiftung for generous financial support.

References

- [1] S. Akbulut, *4-manifolds*. Oxf. Grad. Texts Math. 25, Oxford University Press, Oxford, 2016 MR [3559604](#)
- [2] M. Atiyah, Topological quantum field theories. *Inst. Hautes Études Sci. Publ. Math.* (1988), no. 68, 175–186 (1989) Zbl [0692.53053](#) MR [1001453](#)
- [3] B. Bakalov and A. Kirillov, Jr., *Lectures on tensor categories and modular functors*. Univ. Lecture Ser. 21, American Mathematical Society, Providence, RI, 2001 Zbl [0965.18002](#) MR [1797619](#)
- [4] B. Balsam, *Turaev–Viro theory as an extended TQFT*. ProQuest LLC, Ann Arbor, MI, 2012 MR [3093932](#)

- [5] M. Bärenz and J. Barrett, [Dichromatic state sum models for four-manifolds from pivotal functors](#). *Comm. Math. Phys.* **360** (2018), no. 2, 663–714 Zbl [1414.57017](#) MR [3800795](#)
- [6] M. Barkeshli, P. Bonderson, M. Cheng, and Z. Wang, [Symmetry, defects, and gauging of topological phases](#). 2014, arXiv:[1410.4540v1](#)
- [7] J. W. Barrett, C. Meusburger, and G. Schaumann, [Gray categories with duals and their diagrams](#). 2012, arXiv:[1211.0529](#)
- [8] J. W. Barrett and B. W. Westbury, [Invariants of piecewise-linear 3-manifolds](#). *Trans. Amer. Math. Soc.* **348** (1996), no. 10, 3997–4022 Zbl [0865.57013](#) MR [1357878](#)
- [9] N. Carqueville, I. Runkel, and G. Schaumann, [Orbifolds of \$n\$ -dimensional defect TQFTs](#). *Geom. Topol.* **23** (2019), no. 2, 781–864 Zbl [1441.57030](#) MR [3939053](#)
- [10] L. Crane, L. H. Kauffman, and D. N. Yetter, [State-sum invariants of 4-manifolds](#). *J. Knot Theory Ramifications* **6** (1997), no. 2, 177–234 Zbl [0883.57021](#) MR [1452438](#)
- [11] X. Cui, [Four dimensional topological quantum field theories from \$G\$ -crossed braided categories](#). *Quantum Topol.* **10** (2019), no. 4, 593–676. Zbl [1507.57024](#) MR [4033513](#)
- [12] C. L. Douglas and D. J. Reutter, [Fusion 2-categories and a state-sum invariant for 4-manifolds](#). 2018, arXiv:[1812.11933](#)
- [13] V. Drinfeld, S. Gelaki, D. Nikshych, and V. Ostrik, [On braided fusion categories. I](#). *Selecta Math. (N.S.)* **16** (2010), no. 1, 1–119 Zbl [1201.18005](#) MR [2609644](#)
- [14] P. Etingof, D. Nikshych, and V. Ostrik, [On fusion categories](#). *Ann. of Math. (2)* **162** (2005), no. 2, 581–642 Zbl [1125.16025](#) MR [2183279](#)
- [15] R. E. Gompf and A. I. Stipsicz, [4-manifolds and Kirby calculus](#). Graduate Studies in Mathematics 20, American Mathematical Society, Providence, RI, 1999 Zbl [0933.57020](#) MR [1707327](#)
- [16] J. Harer, A. Kas, and R. Kirby, [Handlebody decompositions of complex surfaces](#). *Grad. Stud. Math.* **62** (1986), no. 350 Zbl [0631.57009](#) MR [849942](#)
- [17] A. Joyal and R. Street, [The geometry of tensor calculus. I](#). *Adv. Math.* **88** (1991), no. 1, 55–112 Zbl [0738.18005](#) MR [1113284](#)
- [18] D. Joyce, [On manifolds with corners](#). In *Advances in geometric analysis*, pp. 225–258, Adv. Lect. Math. (ALM) 21, Int. Press, Somerville, MA, 2012 Zbl [1317.58001](#) MR [3077259](#)
- [19] A. Juhász, [Defining and classifying TQFTs via surgery](#). *Quantum Topol.* **9** (2018), no. 2, 229–321. Zbl [1408.57031](#) MR [3812798](#)
- [20] D. Kasprowski, M. Powell, and P. Teichner, [Four-manifolds up to connected sum with complex projective planes](#). *Amer. J. Math.* **144** (2022), no. 1, 75–118 Zbl [07465467](#) MR [4367415](#)
- [21] R. C. Kirby, [The topology of 4-manifolds](#). Lecture Notes in Math. 1374, Springer, Berlin, 1989 Zbl [0668.57001](#) MR [1001966](#)
- [22] A. A. Kirillov, [String-net model of Turaev–Viro invariants](#). 2011, arXiv:[1106.6033v1](#)
- [23] M. Mackaay, [Finite groups, spherical 2-categories, and 4-manifold invariants](#). *Adv. Math.* **153** (2000), no. 2, 353–390 Zbl [1069.57018](#) MR [1770934](#)
- [24] M. Müger, [On the structure of braided crossed \$G\$ -categories](#). Appendix of V. G. Turaev, *Homotopy quantum field theories*, EMS Tracts Math. 10, European Mathematical Society (EMS), 2010, 221–235

- [25] N. Reshetikhin and V. G. Turaev, [Invariants of 3-manifolds via link polynomials and quantum groups](#). *Invent. Math.* **103** (1991), no. 3, 547–597 Zbl [0725.57007](#) MR [1091619](#)
- [26] J. Roberts, [Skein theory and Turaev–Viro invariants](#). *Topology* **34** (1995), no. 4, 771–787 Zbl [0866.57014](#) MR [1362787](#)
- [27] C. J. Schommer-Pries, *The classification of two-dimensional extended topological field theories*. Ph.D. Thesis. University of California, Berkeley, CA, 2009, arXiv:[1112.1000v1](#)
- [28] M. C. Shum, [Tortile tensor categories](#). *J. Pure Appl. Algebra* **93** (1994), no. 1, 57–110 Zbl [0803.18004](#) MR [1268782](#)
- [29] V. Turaev, Homotopy field theory in dimension 3 and crossed group-categories. 2000, arXiv:[math/0005291v1](#)
- [30] V. G. Turaev, [Quantum invariants of knots and 3-manifolds](#). revised edn., De Gruyter Studies in Mathematics 18, de Gruyter, Berlin, 2010 Zbl [1213.57002](#) MR [2654259](#)
- [31] V. G. Turaev and O. Y. Viro, [State sum invariants of 3-manifolds and quantum 6j-symbols](#). *Topology* **31** (1992), no. 4, 865–902 Zbl [0779.57009](#) MR [1191386](#)
- [32] D. J. Williamson and Z. Wang, [Hamiltonian models for topological phases of matter in three spatial dimensions](#). *Ann. Physics* **377** (2017), 311–344 Zbl [1368.81143](#) MR [3606141](#)

Received 29 October 2018.

Manuel Bärenz

Fakultät für Mathematik, Universität Wien, Oskar-Morgenstern-Platz 1, 1090 Wien, Austria; and eServices GmbH, Blücherstraße 22, 10961 Berlin, Germany; maths@manuelbaerenz.de