View metadata, citation and similar papers at core.ac.uk

brought to you by CORE

provided by INRIA a CCSD electronic archive server

Parameterized complexity of quantum knot invariants

Clément Maria

To cite this version:

Clément Maria. Parameterized complexity of quantum knot invariants. 2020. hal-02429767

HAL Id: hal-02429767

https://hal.archives-ouvertes.fr/hal-02429767

Preprint submitted on 6 Jan 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

PARAMETERIZED COMPLEXITY OF QUANTUM KNOT INVARIANTS

CLÉMENT MARIA

Abstract. We give a general fixed parameter tractable algorithm to compute quantum invariants of links presented by diagrams, whose complexity is singly exponential in the carving-width (or the tree-width) of the diagram.

In particular, we get a $O(N^{\frac{3}{2}\,\mathrm{cw}}\mathrm{poly}(n))$ time algorithm to compute any Reshetikhin-Turaev invariant—derived from a simple Lie algebra $\mathfrak{g}$—of a link presented by a planar diagram with n crossings and carving-width cw, and whose components are coloured with $\mathfrak{g}$-modules of dimension at most N. For example, this includes the Nth-coloured Jones polynomial and the Nth-coloured HOMFLYPT polynomial.

1. Introduction

In geometric topology, testing the topological equivalence of knots (up to isotopy) is a fundamental yet remarkably difficult algorithmic problem.

A main approach is to compare knots by properties depending on their topological types, called *invariants*. Starting with the introduction by Jones [15] in the 1980s of a new polynomial invariant of knots, we have witnessed the birth of a new domain of low dimensional topology called *quantum topology*. From the study of quantum groups [5, 14] in algebra, topologists have designed new families of topological invariants for knots, links, and 3-manifolds, such as the Reshetikhin-Turaev invariants [21]. In practice, these *quantum invariants* have shown outstanding discriminative properties for non-equivalent knots and links, e.g., in the composition of knot census databases [2], and are at the heart of deep mathematical conjectures in the field [7, 8, 16, 20].

Consequently, efficient algorithms to compute quantum invariants are of strong interest. However, even the simplest quantum invariants, such as the Jones polynomial [13], are #P-hard to compute. A successful approach towards practical implementations has been the introduction of *parameterized complexity* to low dimensional topology. Independently, computing the Jones polynomial [18] and the HOMFLYPT polynomial [3] have been shown to admit fixed parameter tractable algorithms in the tree-width of the input link diagrams. Note that similar techniques have been applied to quantum invariants of 3-manifolds, such as the Barrett-Westbury-Turaev-Viro invariants of triangulated 3-manifolds [4, 27]. These algorithms led to significant speed-ups in practice.

Contribution. In this article, we give an algorithm to compute quantum invariants derived from ribbon categories [21, 26], taking into account the carving-width of the input link diagram.

Theorem 1.1. *Fix a strict ribbon category $\mathcal{C}$ of $\mathbb{Z}[q]$-modules, and free modules $V_1, \dots, V_m \in \mathcal{C}$ of dimension bounded by N. The problem:*

> Quantum invariant at $\mathcal{C}, V_1, \dots, V_m$:
> **Input**: *m-components link L, presented by a diagram $D(L)$,*
> **Output**: *quantum invariant $J_L^{\mathcal{C}}(V_1, \dots, V_m)$*

can be solved in $O(\mathrm{poly}(n)N^{\frac{3}{2}\,\mathrm{cw}}) \in O(\mathrm{poly}(n)N^{\frac{3}{2}\sqrt{n}})$ *machine operations, with* $O(N^{\mathrm{cw}} + n)$ *memory words, where* n *and* cw *are respectively the number of crossings and the carving-width of the diagram* $D(L)$*.*

In particular, this implies that, up to some easily computable re-normalisation, computing any Reshetikhin-Turaev invariant derived from a simple Lie algebra $\mathfrak{g}$ is *fixed parameter tractable* (complexity class `FPT`) in the carving-width of the input link diagram. Cases of interests are, in particular, $\mathfrak{g} = \mathfrak{sl}(2, \mathbb{C})$ giving the Nth-coloured Jones polynomials, and $\mathfrak{g} = \mathfrak{sl}(n, \mathbb{C})$ giving the Nth-coloured HOMFLYPT polynomials. This algorithm is:

1 the first fixed parameter tractable algorithm, and—considering $\mathrm{cw} = O(\sqrt{n})$—sub-exponential time algorithm, for quantum invariants of knots stated in such generality (previously known cases were the (uncoloured) Jones polynomial [18], and the (uncoloured) HOMFLYPT polynomial [3]),
2 an exponential improvement over Burton's $2^{O(\mathrm{cw}\log\mathrm{cw})}\mathrm{poly}(n)$ time algorithm for the uncoloured HOMFLYPT polynomial [3], and generally a low exponent ($\frac{3}{2}\,\mathrm{cw}$) singly exponential algorithm for quantum invariants[1].

In Section 2 we recall the definition of quantum invariants derived from ribbon categories, and notions of parameterized complexity. In Section 3 we introduce a high-level parameterized algorithm based on graphical calculus and a tree embedding, then detail in Section 4 the main operation of the algorithm. In Section 5 we develop the implementation of the algorithm in the case of a ribbon category of $\mathcal{R}$-modules, and analyse its arithmetic complexity in Section 6, in the case $\mathcal{R} = \mathbb{Z}[q]$. This last study concludes the proof of the main theorem, and implies more generally that, when the type of invariant is part of the input, computing a quantum invariant is in the complexity class `XP`.

2. Background

We introduce the necessary notions from knot theory, quantum topology, and parameterized complexity.

Tangles and diagrams. A *tangle* is a piecewise linear embedding of a collection of arcs and circles into $\mathbb{R}^2 \times [0, 1]$, such that the arcs' endpoints, called *bases*, belong to the top or bottom boundaries $\mathbb{R}^2 \times \{0\}$ and $\mathbb{R}^2 \times \{1\}$. A tangle intersecting i times $\mathbb{R}^2 \times \{0\}$ and j times $\mathbb{R}^2 \times \{1\}$ is an (i, j)-tangle.

A *link* is a tangle whose connected components are all closed curves (a $(0, 0)$-tangle), and a *knot* is a 1-component link. An *orientation* on a tangle is an orientation of each tangle component. Two tangles are equivalent iff they differ by an ambient isotopy of $\mathbb{R}^2 \times [0, 1]$ maintaining the boundary fixed.

A tangle *diagram* is a projection of the tangle into the plane, induced by a projection of $\mathbb{R}^2 \times [0, 1]$ into $\mathbb{R} \times [0, 1]$, sending $\mathbb{R}^2 \times \{0\}$ and $\mathbb{R}^2 \times \{1\}$ to $\mathbb{R} \times \{0\}$ and $\mathbb{R} \times \{1\}$ respectively. In a tangle diagram, the only multiple points are *crossings*, at which one section of the tangle crosses under or over another one transversally. We consider diagrams with no base points (i.e., link diagrams) as living on the sphere S^2.

Component orientations are pictured with arrow heads, and a $k \in \mathbb{Z}$ framing is pictured by k *positive twists* if $k > 0$, and k *negative twists* is $k < 0$. See Figure 1.

We refer to [17] for more details on knot theory.

[1]Note that previous algorithms [18] are expressed in terms of tree-width, which is proportional but not equal to the carving-width, in consequence exponents are not directly comparable.

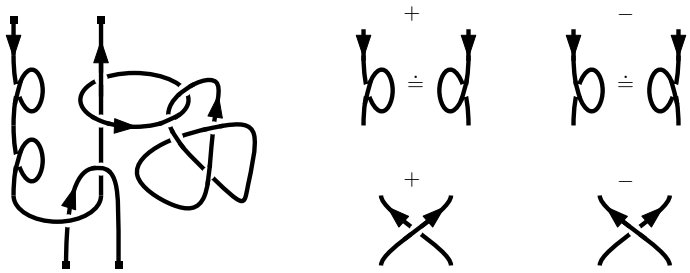

FIGURE 1. Left: Diagram of a 4-components oriented framed tangle, whose top left component has framing +2. Right: Positive/negative twists and crossings. The $\doteq$ symbol is an equivalence of diagrams.

Ribbon categories and quantum invariants. We refer to Turaev's monograph [26] for the categorical formulation of quantum invariants. We only introduce the necessary notions.

Intuitively, a strict ribbon category is an abstraction of the category of modules over a commutative ring, with their usual tensor product. Some morphisms—called *braidings*, *twists*, *evaluations* and *co-evaluations*—are distinguished in order to establish a connection between topology (tangles and knots) and algebra (morphisms between objects), via graphical calculus.

More precisely, a *strict ribbon category* $\mathcal{C}$ is a category with a unit object $\mathbb{1}$ and which is equipped, for any objects U, V, U', V' and morphisms $f\colon V \to V'$, $g\colon U \to U'$, with:

(a) an associative tensor product assigning to U and V an object $U \otimes V$, and to f and g a morphism $f \otimes g\colon U \otimes V \to U' \otimes V'$,
(b) a natural braiding isomorphism $c_{U,V}\colon U \otimes V \to V \otimes U$ with inverse $c_{U,V}^{-1}$,
(c) a duality associating to any V a dual object V^*, together with *co-evaluation* morphisms $b_V\colon \mathbb{1} \to V \otimes V^*$ and *evaluation* morphisms $d_V\colon V^* \otimes V \to \mathbb{1}$,
(d) a natural twist isomorphism $\theta_V\colon V \to V$ with inverse θ_V^{-1},
(e) and where $\mathrm{Hom}_{\mathcal{C}}(\mathbb{1}, \mathbb{1})$ has the structure of a commutative ring $\mathcal{R}$.

By convention, the "tensor product of zero objects" is equal to $\mathbb{1}$. In a strict ribbon category, these objects and morphisms satisfy additional compatibility constraints, that are necessary to state Theorem 2.1 below.

For example, the category of modules over a commutative ring $\mathcal{R}$ with standard tensor product, and equipped with the trivial braiding $u \otimes v \mapsto v \otimes u$, forms a strict ribbon category. In this case, the ring $\mathcal{R}$, seen as a module over itself, is the unit object $\mathbb{1}$, and any morphism $\mathcal{R} \to \mathcal{R}$ is a multiplication by a scalar $\tau \in \mathcal{R}$. Hence $\mathrm{Hom}_{\mathcal{C}}(\mathbb{1}, \mathbb{1})$ is isomorphic to the commutative ring $\mathcal{R}$ itself. For invariants derived from quantum groups, we mainly focus on the category of $\mathcal{R}$-modules, generally free of finite dimension but with more complex braidings than the trivial ones. The ring $\mathcal{R}$ is $\mathbb{Z}[q]$ (or $\mathbb{Z}[q, q^{-1}]$), the ring of one-variable polynomials with integer coefficients. Morphisms between free modules are represented by matrices with $\mathcal{R}$-coefficients.

Graphical calculus and coloured tangles. Fix a strict ribbon category $\mathcal{C}$. A *colouring* of a link L, with m ordered components $L_1, \ldots, L_m$, is an assignment of an object $V_i \in \mathcal{C}$, $1 \leq i \leq m$, to every component L_i of L.

A link diagram is considered in *standard form* if it can be decomposed into the following pieces, described in Figure 2: (i) vertical strands, (vi) & (vii) positive and negative crossings, (viii) & (ix) positive and negative right twists, and (x) & (xi) caps and cups. See Figure 3 for a Hopf link in standard position. Any link (or tangle) diagram can be moved into standard form.

Rules (i) to (xi) of Figure 2 gives the conversion from coloured tangle to $\mathcal{C}$-morphism, called *Penrose functor*. Specifically, given a coloured link diagram $D(L)$ in standard form, the Penrose functor turns the diagram into a morphism, following the rules:

(o): A morphism $f\colon U \to V$ in $\mathcal{C}$ is represented graphically by a box, aligned with x- and $y-$axis, called *coupon*, with incoming vertical V-coloured strands (top) and outgoing vertical U-coloured strand (bottom),

(i): reversing a component orientation changes a colour V to its dual V^*,

(ii): two parallel strands coloured U and V are equivalent to a single strand coloured $U \otimes V$,

(iii): a vertical strand coloured V is equivalent to the identity morphism id_V,

(iv): a morphism g above another one f is equivalent to there composition $g \circ f$,

(v): two morphisms h_1 and h_2 side by side are equivalent to their tensor product $h_1 \otimes h_2$,

(vi) & (vii): a positive crossing is equivalent to a braiding morphism, a negative crossing is equivalent to the inverse of the braiding morphism,

(viii) & (ix): positive and negative twists are equivalent to the twist morphism and its inverse respectively,

(x) & (xi): caps and cups are equivalent to evaluation and co-evaluation respectively.

(xii): the dual morphism $f^*\colon V^* \to U^*$ of a morphism $f\colon U \to V$ is given by the graphical equation (xii) or, equivalently, by:

$$f^* = (d_V \otimes \mathrm{id}_{U^*}) \circ (\mathrm{id}_{V^*} \otimes f \otimes \mathrm{id}_{U^*}) \circ (\mathrm{id}_{V^*} \otimes b_U).$$

The morphisms are applied to the objects colouring the entering and leaving strands. Figure 2 gives the morphism associated to the Hopf link coloured with objects U and V.

Consequently, for a category $\mathcal{C}$, the Penrose functor associates to any coloured link a morphism $\mathbb{1} \to \mathbb{1}$. More generally, it associates to a coloured (i, j)-tangle a morphism $U_1 \otimes \ldots \otimes U_i \to V_1 \otimes \ldots \otimes V_j$, for the $\mathcal{C}$-objects U_ks and V_ℓs colouring the bottom and top bases respectively.

If the ordered components of a link L are coloured $V_1, \ldots, V_m$, this morphism is written:

$$J_L^{\mathcal{C}}(V_1, \ldots, V_m) \in \mathrm{Hom}_{\mathcal{C}}(\mathbb{1}, \mathbb{1}).$$

Strict ribbon category produce topological invariants, called *quantum invariants*:

Theorem 2.1 ([21, 26])**.** *Let $D(L)$ be a diagram of an m-components link L on S^2, and let $\mathcal{C}$ be a strict ribbon category. Let $V_1, \ldots, V_m$ be a colouring of the components of L. The quantity $J_L^{\mathcal{C}}(V_1, ..., V_m)$ produced by the Penrose functor is invariant by ambient isotopy of S^2 and Reidemeister moves on $D(L)$. It is consequently a topological invariant of the coloured link L.*

When $\mathcal{C}$ is the category of $\mathcal{R}$-modules, $J_L^{\mathcal{C}}(V_1, ..., V_k) \in \mathrm{Hom}(\mathbb{1}, \mathbb{1}) \cong \mathcal{R}$ is identified to a scalar in $\mathcal{R}$.

Graph parameters. The *carving-width*, also known as *congestion*, is a graph parameter introduced by Seymour and Thomas [25].

Definition 2.2. Let $\mathcal{G} = (V, E)$ be a graph on n vertices, with loops and multiple edges. Let $\mathcal{T}$ be an unrooted binary tree, with all internal nodes of degree 3, and with n leaves.

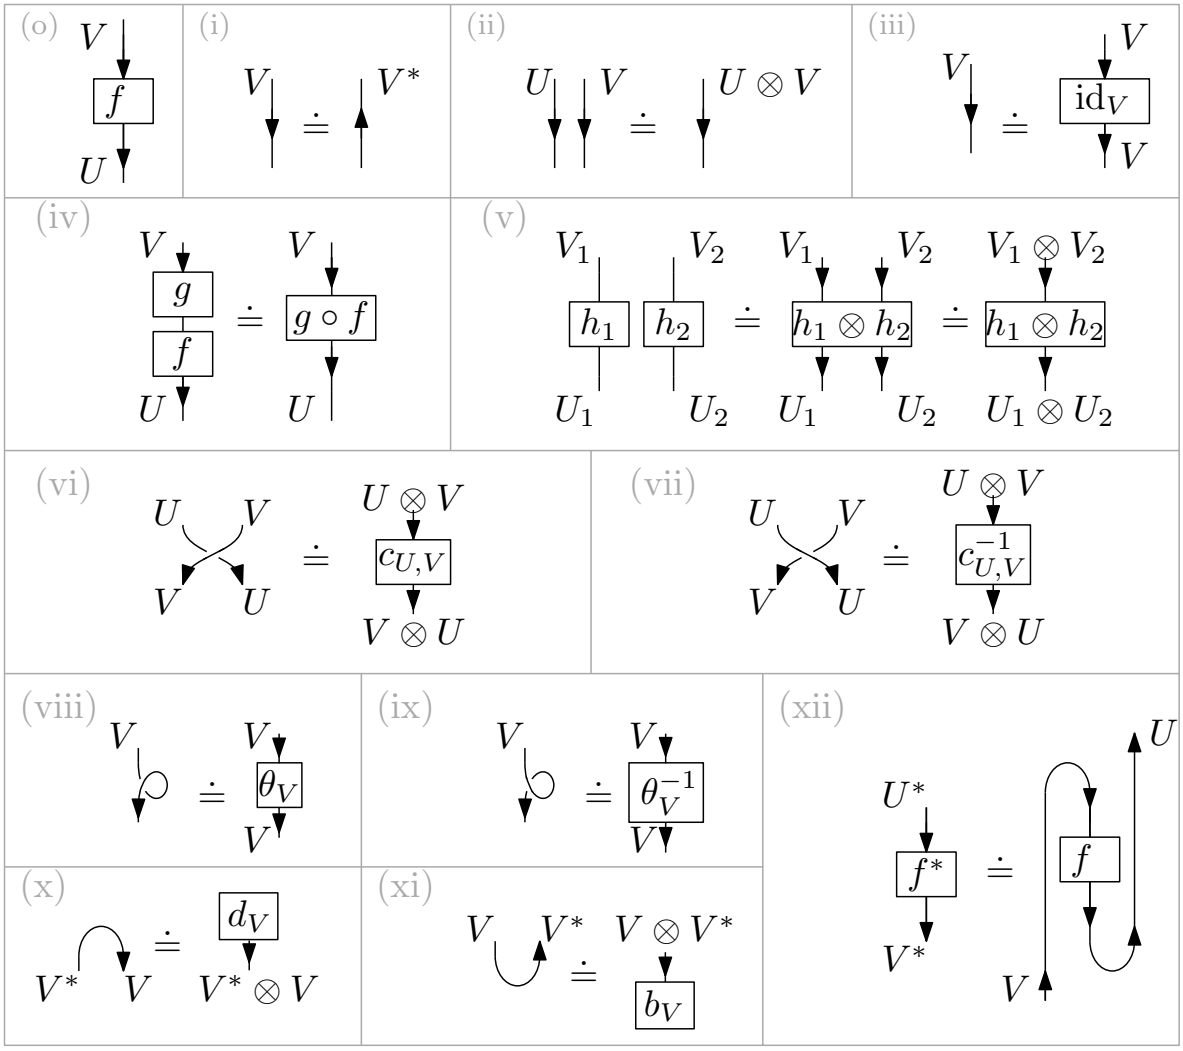

FIGURE 2. Graphical calculus induced by Penrose functor.

An *embedding* ϕ of $\mathcal{G}$ into $\mathcal{T}$ is a bijective mapping between the nodes of $\mathcal{G}$ and the leaves of $\mathcal{T}$. Every edge e of $\mathcal{T}$ induces a partition of the vertices of $\mathcal{G}$ into two sets, $V = U_e \sqcup V_e$, inherited from the partition of $\mathcal{T} \setminus e$ into two trees. Let $\mathrm{w}(e)$ denote the number of edges in $\mathcal{G}$ between U_e and V_e, called the weight of e.

The *congestion* of an embedding $(\mathcal{T}, \phi)$ is the maximal weight of a tree edge:

$$\mathrm{cng}(\mathcal{T}, \phi) = \max_{e \text{ edge of } \mathcal{T}} \mathrm{w}(e),$$

The *carving-width* $\mathrm{cng}(\mathcal{G})$ of a graph $\mathcal{G}$ is the minimal congestion over all its embeddings into binary trees. The *carving-width* $\mathrm{cng}(D(L))$ *of a link diagram* $D(L)$ is the carving-width of the 4-valent planar graph it realises. The carving-width $\mathrm{cng}(L)$ of a link L is the minimal carving-width of any of its diagrams.

The carving-width of a graph is closely related to its *tree-width* [22], which plays a major role in combinatorial algorithms.

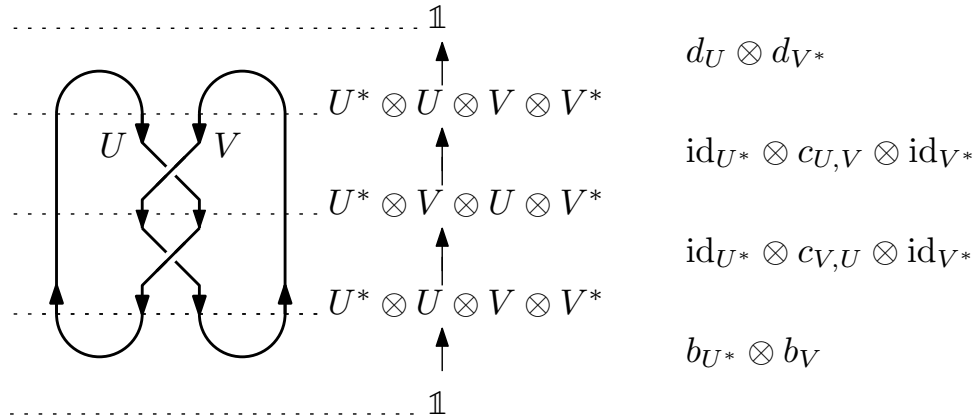

FIGURE 3. Application of Penrose functor to the Hopf link coloured by objects U and V from a strict ribbon category, leading to a $\mathbb{1} \to \mathbb{1}$ morphism by composition.

Theorem 2.3 (Theorem 1 of [1])**.** *Let $\mathcal{G}$ be a graph of maximal degree δ. Then,*

$$\frac{2}{3}(\mathrm{tw}(\mathcal{G})+1) \leq \mathrm{cng}(\mathcal{G}) \leq \delta(\mathrm{tw}(\mathcal{G})+1). \qquad \textit{For tangle diagrams } \delta \leq 4.$$

Carving-width has several advantages over tree-width, and has been successfully used in low dimensional topology [11, 12, 19, 24].

First, optimal tree embeddings of planar graphs can be realised topologically, as stated below. A *bridge* in a connected graph G is an edge of G whose removal splits G into more than one connected component. A tree embedding $(\mathcal{T}, \phi)$ of G is *bond* if the two vertex sets U_e and V_e from the cut associated to an edge e of $\mathcal{T}$ induce connected sub-graphs in G.

Theorem 2.4 ([25, Theorem 5.1])**.** *Let G be a simple connected bridgeless graph with more than two vertices. If G has carving-width* cw *then there exists a bond tree embedding of G of congestion* cw*.*

Up to a subdivision of multiple edges, which does not increase carving-width, a link diagram can be made simple, as a graph. Being 4-valent, it is bridgeless, and, if connected, it consequently admits a bond tree embedding of minimal congestion. We interpret a bond tree embedding of a planar graph (on the sphere S^2) as a collection of disjoint Jordan curves $\lambda_e \subset S^2$, one for each edge e of $\mathcal{T}$, realising the cut $U_e \sqcup V_e$ [24].

For planar graphs, a bond tree embedding of minimal congestion can be computed in polynomial time [10, 25]. By the planar separator theorem, the carving-width of a planar graph with n-vertices is in $O(\sqrt{n})$.

3. Fixed parameter tractable algorithm via graphical calculus

Let $\mathcal{C}$ be a strict ribbon category, and let L be an oriented link with m components $L_1, \ldots, L_m$. Let $D(L)$ be an oriented link diagram of L, where each link component L_i is coloured by an object V_i from the category $\mathcal{C}$, such that the Penrose functor gives an isotopy invariant of L associated to its colouring, as described in Theorem 2.1.

It follows from the definition of Penrose functor that the quantum invariant of a separable link $L \cup L'$ is the product of the invariants of L and L', such that they can be computed separately. W.l.o.g. we assume that the diagram $D(L)$ is connected as a graph, and has at least 2 crossings, not all twists.

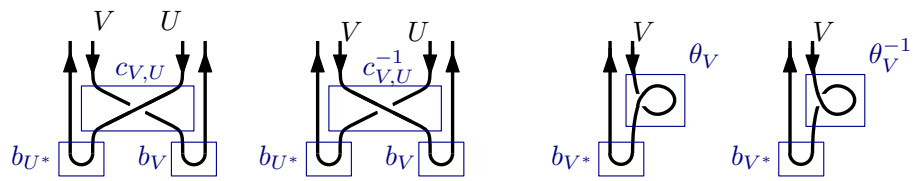

FIGURE 4. The four tangles and associated morphisms at the tree leaves. From left to right: Equations (3.1), (3.2), (3.3), and (3.4). The marked bullet point is on the left of each diagram, and is selected such that only these four morphisms are encountered.

3.1. **Tree embedding of link diagrams.** Let $(\mathcal{T}, \phi)$ be a bond tree embedding of the planar graph of $D(L)$, and root it by subdividing an arbitrary tree edge, picking the centre as the root. All edges of $\mathcal{T}$ have now a parent and child endpoint. By convention, we add a "half-edge" on top of the tree, having the root as child. Every inner node in $\mathcal{T}$ has consequently degree 3, with two edges "going down", and one edge "going up".

Let e be an edge of $\mathcal{T}$ with child node x, and X the set of crossings mapped to the leaves of the subtree $\mathcal{T}_x$ rooted at x. According to Theorem 2.4, there exists a Jordan curve λ_e separating X from the rest of the diagram. The diagram being on the sphere, we draw the tangle "inside" the Jordan curve when we represent it on the plane.

To an edge e corresponds a $(0, \mathrm{w}(e))$-tangle T, spanned by the crossings X and contained "inside" λ_e. We mark an arbitrary but fixed "bullet" point on λ_e and order the bases of T counter-clockwise from that point. We get a $(0, \mathrm{w}(e))$-tangle by isotopically sliding all bases to the top boundary, such that the first base in the bullet ordering is rightmost on the top boundary. See Figure 4 for examples of $(0, \mathrm{w}(e))$-tangles at the tree leaves, and Figure 5 (Left) for bases ordered by bullet ordering.

In the process of the algorithm below, bullet orderings are assigned on the fly.

3.2. **Tree traversal algorithm.** Let $D(L)$ be coloured by objects of the category $\mathcal{C}$. To every edge e of weight $\mathrm{w}(e)$ in $\mathcal{T}$, the Penrose functor assigns a $\mathcal{C}$-morphism $f_e \colon \mathbb{1} \to V_1 \otimes \ldots \otimes V_{\mathrm{w}(e)}$ to the associated tangle, where $V_1, \ldots, V_{\mathrm{w}(e)}$ are the colours of the strands intersecting the Jordan curve λ_e.

The morphism associated to the half-edge at the root is a $\mathbb{1} \to \mathbb{1}$ morphism, because the corresponding Jordan curve does not intersect the link diagram. This morphisms gives the invariant $J_L^{\mathcal{C}} \in \mathcal{R}$ of Theorem 2.1. All edge morphisms are computed recursively following a depth first traversal of $\mathcal{T}$. We describe the base morphisms assigned to the edges whose child node is a leaf, and we describe an algorithm for inner edges in the next section.

3.3. **Morphisms at the leaves.** Up to reorientation of the strands, which algebraically consists of dualising colours, we can restrict to four base morphisms:

$$(\mathrm{id}_{U^*} \otimes c_{V,U} \otimes \mathrm{id}_{V^*}) \circ (b_{U^*} \otimes b_V) \tag{3.1}$$

$$(\mathrm{id}_{U^*} \otimes c^{-1}_{V,U} \otimes \mathrm{id}_{V^*}) \circ (b_{U^*} \otimes b_V) \tag{3.2}$$

$$(\mathrm{id}_{V^*} \otimes \theta_V) \circ b_{V^*} \tag{3.3}$$

$$(\mathrm{id}_{V^*} \otimes \theta_V^{-1}) \circ b_{V^*} \tag{3.4}$$

They correspond graphically to the diagrams in Figure 4, where the bullet ordering is chosen to restrict to these four cases.

3.4. **Merging morphisms at tree nodes.** Every inner node x of $\mathcal{T}$ is the parent node of two edges e_1 and e_2, and the child of an edge e. Given the morphisms f_{e_1} and f_{e_2} for edges e_1 and e_2 respectively, we construct the morphism f_e for edge e.

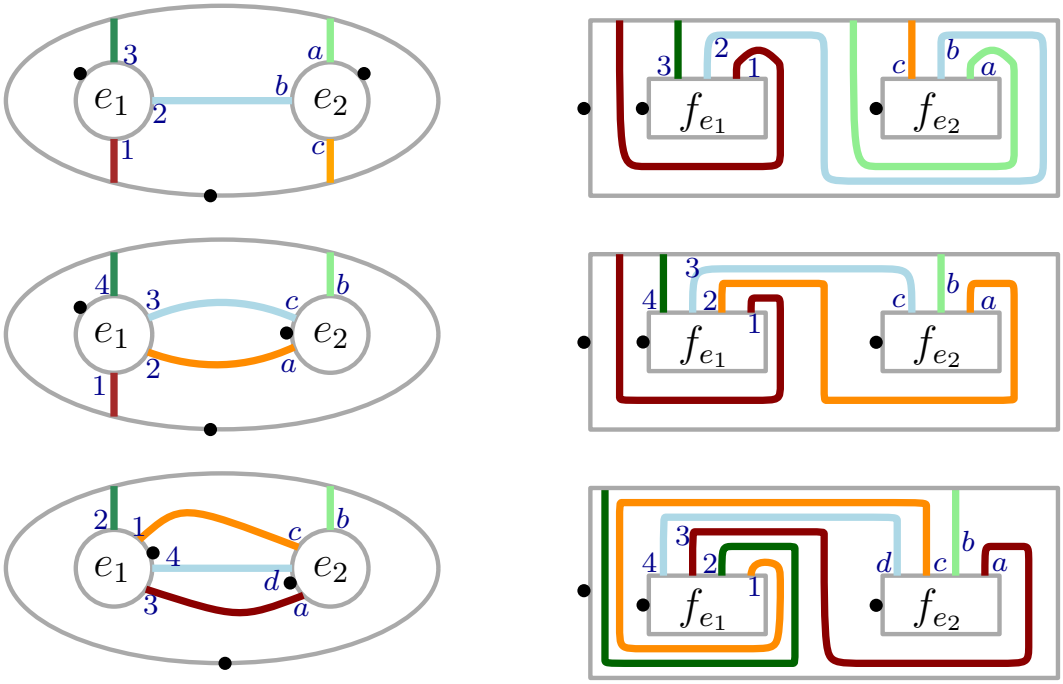

FIGURE 5. Merging two sub-trees. Left: Planar embeddings of the diagram with Jordan curves λ_{e_1}, λ_{e_2} (inner circles) and λ_e (outer circle), depending on the position of the bullets for λ_{e_1} and λ_{e_2}. The bold lines connecting the Jordan curves represent multiple parallel strands connecting the corresponding tangles. Right: Coupons for f_{e_1}, f_{e_2} and f_e (outer coupon) obtained after plane isotopy. The bullet for λ_e is selected so as to restrict to these three cases.

First, note that the bullet ordering of the strands intersecting λ_{e_1} and λ_{e_2} leads to three configurations when representing morphisms f_{e_1} and f_{e_2} with coupons ; see Figure 5 where thick lines represent sets of parallel tangle strands. By hypothesis, morphisms on tree edges have domain $\mathbb{1}$. The coupons for f_{e_1}, f_{e_2}, and f_e (the outer coupon) are obtained by a plane isotopy forcing the strands to intersect coupons on their top side, and putting bullets on the coupons' left sides. The bullet of the outer coupon f_e is selected so as to restrict to the three configurations of Figure 5.

4. Factorisation of morphisms at tree nodes

Given the morphisms f_{e_1} and f_{e_2} in Figure 5, we describe graphically a factorisation scheme to obtain the morphism f_e.

4.1. **Sliding and canonical form.** The canonical form for morphisms to be merged is depicted in the top left corner of Figure 7. It consists of two side-by-side morphisms g_1 and g_2, bridged by parallel strands coloured $U_1, \dots U_k$. All other strands go vertically.

Given morphisms f_{e_1} and f_{e_2} in Figure 5, we obtain a canonical form by sliding strands, wrapping clockwise around the coupons, under the coupons. For example, in the top right case of Figure 5, we slide strand 1 under the f_{e_1}-coupon, and strands a and b under the f_{e_2}-coupons.

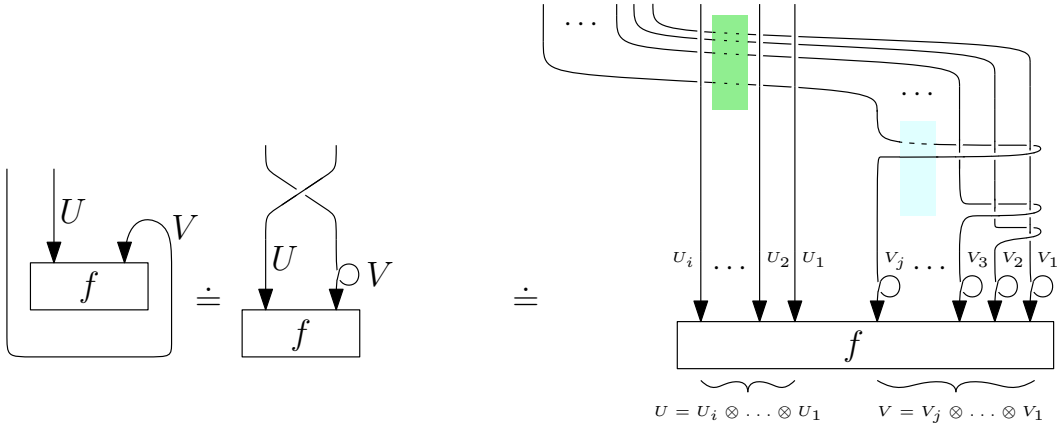

FIGURE 6. Sliding of a $V = V_j \otimes \ldots \otimes V_1$-coloured strand under an f-coupon by underlying knot isotopy. The operation composes f with a consecutive sequence of j twists θ_{V_ℓ}, of $j(j-1)$ crossings $c^{\pm}_{V_i,V_j}$, and ij crossings $c^{\pm}_{U_k,V_\ell}$.

The details of the operation are depicted in Figure 6, where the V-strand wraps clockwise around the f-coupon, and f is a $\mathbb{1} \to U \otimes V$ morphism. Sliding the V-strand under the coupon by tangle isotopy produces a positive twist θ_V and a positive crossing $c_{V,U}$.

Decomposing further in Figure 6, let $U = U_i \otimes \ldots \otimes U_1$ be the tensor product of the colours of i parallel strands, and $V = V_j \otimes \ldots \otimes V_1$ the tensor product of j parallel strands wrapping clockwise around the f coupon. As depicted in the figure, sliding the j strands under f induces

- a twist θ_{V_ℓ} on each of the V_ℓ-coloured strands, $1 \leq \ell \leq j$,
- a sequence of $j(j-1)$ positive and negative crossings of type $c^{\pm}_{V_\ell,V_k}$, followed by
- a sequence of ij positive crossings of type c_{V_ℓ,U_k}.

We obtain the morphisms g_1, g_2 of the canonical form (Figure 7) by factorising the morphisms f_{e_1} and f_{e_2} with these sequences of twists and crossings, after the sliding operation.

4.2. **Factorisation of the canonical form.** Figure 7 pictures two factorisation schemes for side-by-side morphisms g_1 and g_2 in canonical form, bridged by k parallel strands coloured $U_1, \ldots, U_k$. Denote by cw the carving-width of the link diagram, and assume the tree embedding $(\mathcal{T}, \phi)$ has width cw. We distinguish between two cases:
Small bridge. For k smaller than half the carving-width (Figure 7, Left), we consider first the morphism $d_{U_1 \otimes \ldots \otimes U_k}$ induced by the composition of the evaluation morphisms d_{U_ℓ}, $\ell = k \ldots 1$. More precisely, the morphism $d_{U_1 \otimes \ldots \otimes U_k} \colon U_1 \otimes \ldots \otimes U_k \otimes U_k^* \otimes \ldots \otimes U_1^* \to \mathbb{1}$, is obtained by composing the evaluation morphisms from bottom up:

$$\begin{aligned} d_{U_1 \otimes \ldots \otimes U_k} \quad &: \quad U_1 \otimes \ldots \otimes U_k \otimes U_k^* \otimes \ldots \otimes U_1^* \to \mathbb{1}, \\ &= \quad \textstyle\prod_{\ell=k}^{1} \left(\mathrm{id}_{U_1 \otimes \ldots \otimes U_{\ell-1}} \otimes d_{U_\ell^*} \otimes \mathrm{id}_{U_{\ell-1}^* \otimes \ldots \otimes U_1^*} \right) \end{aligned} \tag{4.1}$$

where $\ell = k$ is the rightmost term of the composition.

The (partial) composition of $d_{U_1 \otimes \ldots \otimes U_k}$ with g_2 through $U_k^* \otimes \ldots \otimes U_1^*$ gives the morphism h:

$$\begin{aligned} h \quad &: \quad U_1 \otimes \ldots \otimes U_k \to W_1 \otimes \ldots \otimes W_j, \\ &= \quad \left(d_{U_1 \otimes \ldots \otimes U_k} \otimes \mathrm{id}_{W_1 \otimes \ldots \otimes W_j} \right) \circ \left(\mathrm{id}_{U_1 \otimes \ldots \otimes U_k} \otimes g_2 \right). \end{aligned} \tag{4.2}$$

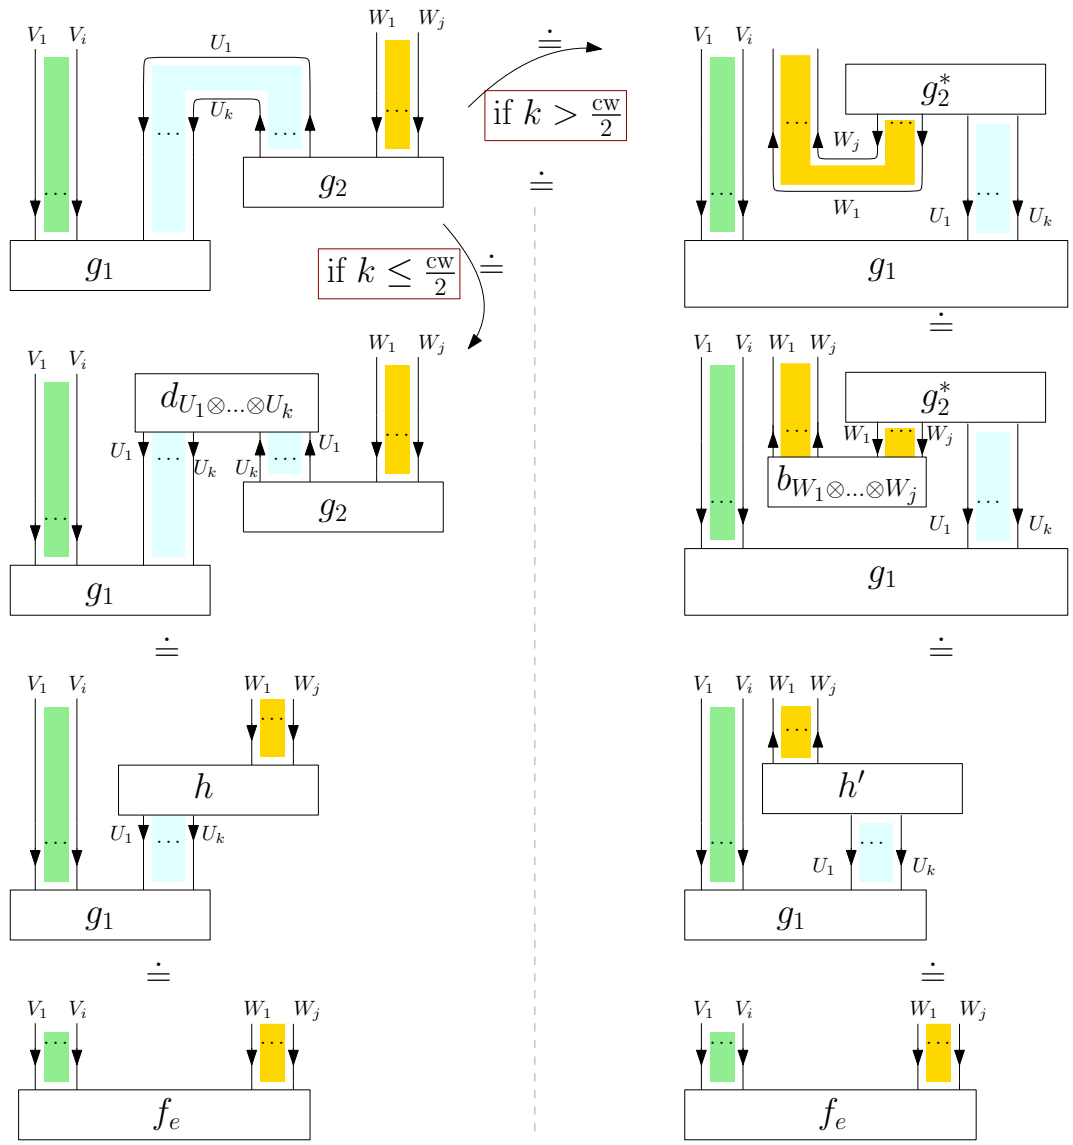

FIGURE 7. Merging of two coupons in a canonical form (top left) along k strands coloured $U_1, \ldots, U_k$. The factorisation scheme differs whether $k \leq \mathrm{cw}/2$ (left column) or $k > \mathrm{cw}/2$ (right column). The top right equivalence comes from the equality in Figure 8.

Finally, the morphism f_e obtained from the merging of f_{e_1} and f_{e_2} is given by the (partial) composition of g_1 and h, through $U_1 \otimes \ldots \otimes U_k$. Precisely,

$$\begin{aligned} f_e &: \mathbb{1} \to V_1 \otimes \ldots \otimes V_i \otimes W_1 \otimes \ldots \otimes W_j, \\ &= (\mathrm{id}_{V_1 \otimes \ldots \otimes V_i} \otimes h) \circ g_1. \end{aligned} \tag{4.3}$$

By construction, these operations give the morphism f_e induced by the Penrose functor on the coloured tangle associated to the subtree of $\mathcal{T}$ rooted at the child node of edge e.

Large bridge. The case k strictly larger than half the carving-width starts by flipping upside-down coupon g_2. Precisely, this operation is depicted in Figure 8. Starting with a morphism

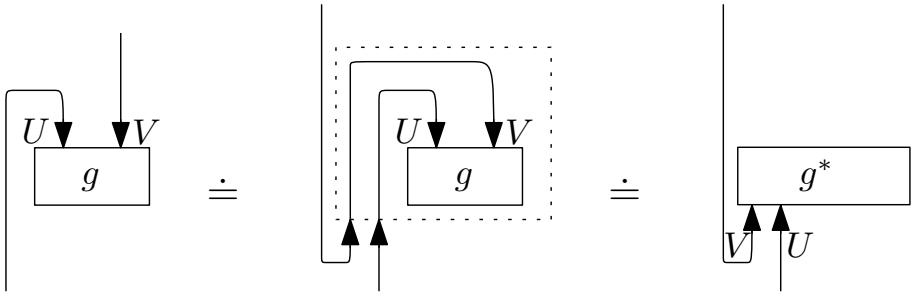

FIGURE 8. Planar isotopy, then factorisation with g^*, the dual morphism to g.

g, it consists of a planar isotopy to produce g^*, the dual morphism to g. In the case where the category $\mathcal{C}$ satisfies the hypothesis of Theorem 2.1, Figure 8, depicting an isotopy, proves the equality:

$$(d_U \otimes \mathrm{id}_V) \circ (\mathrm{id}_{U^*} \otimes g) = (\mathrm{id}_V \otimes g^*) \circ (b_V \otimes \mathrm{id}_{U^*}) .$$

Applied to the canonical form on g_1 and g_2 (Figure 7, Top) the operation gives the composition of morphisms, involving g_1 and g_2^*, in the top right corner of Figure 7.

The following compositions are similar to the case of a small bridge. Morphism $b_{W_1 \otimes \ldots \otimes W_j}$ describes the composition of the co-evaluation morphisms for $W_1, \ldots, W_j$, i.e.,

$$\begin{aligned} b_{W_1 \otimes \ldots \otimes W_j} &: \mathbb{1} \to W_1^* \otimes \ldots \otimes W_j^* \otimes W_j \otimes \ldots \otimes W_1, \\ &= \textstyle\prod_{\ell=1}^{j} \left(\mathrm{id}_{W_1^* \otimes \ldots \otimes W_{\ell-1}^*} \otimes b_{V_\ell} \otimes \mathrm{id}_{W_{\ell-1} \otimes \ldots \otimes W_1} \right) . \end{aligned} \tag{4.4}$$

where $\ell = 1$ is the rightmost term of the composition.

The morphism h' is obtained by (partial) composition of $b_{W_1 \otimes \ldots \otimes W_j}$ and g_2^*:

$$h' = \left(\mathrm{id}_{W_1^* \otimes \ldots \otimes W_j^*} \otimes g_2^* \right) \circ \left(b_{W_1 \otimes \ldots \otimes W_j} \otimes \mathrm{id}_{U_1 \otimes \ldots \otimes U_k} \right) , \tag{4.5}$$

and f_e is obtained by (partial) composition of g_1 and h':

$$f_e = \left(\mathrm{id}_{V_1 \otimes \ldots \otimes V_i} \otimes h' \right) \circ g_1 . \tag{4.6}$$

Correctness. The correctness of the algorithm follows directly from Theorem 2.1, noting that the algorithm consists of an isotopy of the link, realisable by isotopies of the sphere on which the diagram is drawn, and Reidemeister moves.

5. Algebraic implementation and complexity

For the implementation of the algorithm, we assume that the objects in the category $\mathcal{C}$ are finite dimensional free $\mathcal{R}$-modules, for a commutative ring with unity $\mathcal{R}$ and usual tensor product. Denote the dimension of every link component colour V_i by $N_i := \dim V_i$, and let $N := \max_i \{\dim V_i\}$. Fixing a basis for every V_i, all morphisms in $\mathcal{C}$—in particular the distinguished braiding, evaluation and co-evaluation, and twist morphisms—are represented by matrices with $\mathcal{R}$ coefficients.

This model is general, and contains all quantum invariants derived from simple Lie algebras.

5.1. **Elementary compositions.** We consider the seven elementary compositions of morphisms depicted in Figure 9. They respectively represent the composition with (1) a single braiding, (2) a single twist, (3) a single co-evaluation, (4) a single evaluation. Cases (5), (6), and (7) represent three types of partial compositions of the morphisms f and g. We describe algorithms to perform these compositions on matrices.

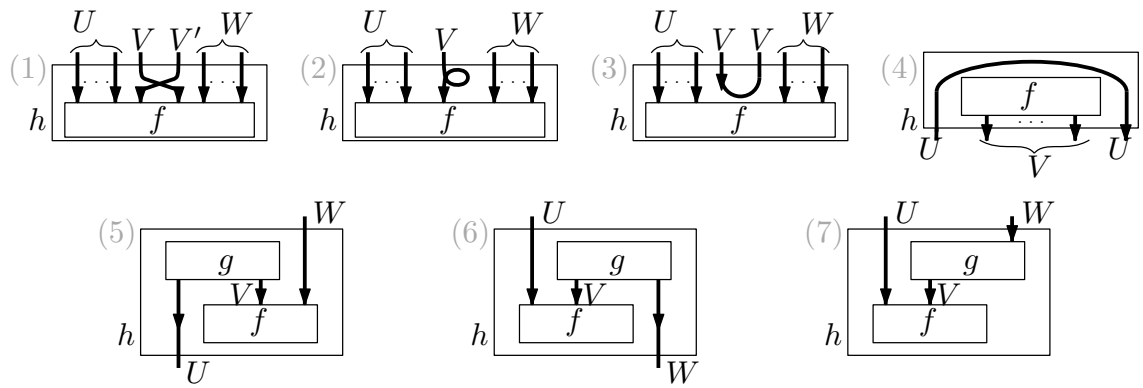

FIGURE 9. Graphical representation of the seven elementary compositions of morphisms.

Lemma 5.1. *Consider the elementary morphism compositions in Figure 9 (1), (2), (3), and (4). Let U, V, V', W be finite dimensional free $\mathcal{R}$-modules, with $\dim U = a$, $\dim V = b$, $\dim V' = b'$, and $\dim W = c$. Then, given the matrices for morphisms f, $\theta_V^{\pm}$, $c_{V,V'}^{\pm}$, b_V, and d_U, we can compute the matrix for morphism h in:*

- *$O(a(bb')^2c)$ arithmetic operations in $\mathcal{R}$ for (1),*
- *$O(ab^2c)$ arithmetic operations for (2) and (3), and*
- *$O(a^2b)$ arithmetic operations for (4).*

The memory complexity of the operation does not exceed the size of the output, which is a row or column vector h containing scalars from $\mathcal{R}$.

Proof. **Figure 9 (1), (2), and (3).** All three cases consist of the matrix-vector product $h = (\mathrm{id}_U \otimes M \otimes \mathrm{id}_W) \cdot f$, where M is respectively the $(bb' \times bb')$-matrix $c_{V,V'}^{\pm}$, the $(b \times b)$-matrix $\theta_V^{\pm}$, and the $(b^2 \times 1)$-matrix b_V.

Consider M to be an $(m \times m')$-matrix, with coefficients $(M_{i,j})_{1 \le i \le m, 1 \le j \le m'}$. Matrix $(\mathrm{id}_U \otimes M \otimes \mathrm{id}_W)$ has at most m' non-zero coefficients per row. We get the formula for the ith entry of h:

$$h_{i,1} = \sum_{k=1 \dots m'} M_{\beta+1,k} \cdot f_{\alpha c m' + \gamma + (k-1)c, 1},$$

where i is uniquely written as $i = \alpha \cdot cm + \beta \cdot c + \gamma$, with $0 \le \alpha \le a-1$, $0 \le \beta \le m-1$, and $1 \le \gamma \le c$. Computing h requires $O(m'|f|)$ arithmetic operation in $\mathcal{R}$, where $|f|$ is the length of vector f, storing $O(|f|)$ scalars from $\mathcal{R}$.

Figure 9 (4). With a similar approach, we get for any j, $1 \le j \le a^2b$:

$$h_{1,j} = (d_U)_{1,\alpha a + \gamma} \cdot f_{\beta+1,1},$$

where j is uniquely written as $j = \alpha \cdot ab + \beta \cdot a + \gamma$, with $0 \le \alpha \le a-1$, $0 \le \beta \le b-1$, and $1 \le \gamma \le a$. The algorithm has complexity $O(a^2b)$ and memory usage $O(a^2b)$. □

Lemma 5.2. *Consider the elementary morphism compositions in Figure 9 (5), (6), and (7). Let U, V, W be finite dimensional free $\mathcal{R}$-modules, with $\dim U = a$, $\dim V = b$, and $\dim W = c$. Then, given the matrices for morphisms f and g, we can compute the matrix for morphism h in $O(abc)$ arithmetic operations in $\mathcal{R}$, and memory complexity $O(ab + bc + ac)$ times the size of a scalar in $\mathcal{R}$.*

Proof. **Figure 9 (5).** Morphism f is a $bc \times 1$ matrix, and morphism g is a $1 \times ab$ matrix. Define $h = (g \otimes \mathrm{id}_W)(\mathrm{id}_U \otimes f)$. Morphism h is a $c \times a$ matrix.

Studying the shape of matrices $(\mathrm{id}_U \otimes f)$ and $(g \otimes \mathrm{id}_W)$, it appears that every one of the $c \times a$ coefficients of the product $h = (g \otimes \mathrm{id}_W)(\mathrm{id}_U \otimes f)$ is a sum of $O(b)$ terms. Precisely, an explicit computation gives us, for any i, j, $1 \leq i \leq c$, $1 \leq j \leq a$:

$$h_{i,j} = \sum_{k=0\ldots b-1} g_{1,(j-1)b+k+1} \cdot f_{kc+i,1}.$$

Morphism h is a $c \times a$ matrix, and each of its coefficients can be computed in $O(b)$ arithmetic operations in $\mathcal{R}$, leading to the $O(abc)$ time complexity. The memory consumption is the sum of the sizes of the input matrices f and g, and the output matrix h.

Figure 9 (6). With a similar approach, for any i, j, $1 \leq i \leq a$, $1 \leq j \leq c$:

$$h_{i,j} = \sum_{k=0\ldots b-1} g_{1,kc+j} \cdot f_{(i-1)b+k+1,1}.$$

Figure 9 (7). With a similar approach, for any i, $1 \leq i \leq ac$:

$$h_{i,1} = \sum_{k=1\ldots b} g_{\beta,k} f_{\alpha b+k,1},$$

where we write $i = \alpha c + \beta$, for $0 \leq \alpha \leq a-1$ and $1 \leq \beta \leq c$. □

5.2. **Implementation of the algorithm.** We implement the algorithm described in Sections 3 and 4 using the elementary composition of Figure 9. Define N a bound on the dimension of the different modules U_i, V_j, W_k colouring the components of the link.
Leaf morphisms. The leaf morphisms described in Equations (3.1-3.1) and Figure 4 are implemented using elementary compositions (1) and (2). By Lemma 5.1, the complexity is at most $O(N^6)$ arithmetic operations in $\mathcal{R}$.
Sliding under a coupon. The sliding operation as presented in Figure 6 composes a morphism f with a sequence of twist and braiding morphisms. Precisely, let h denote the entire morphism in Figure 6. Starting from the $(O(N^{i+j}) \times 1)$ matrix f, it is computed iteratively applying j times elementary composition (2) for the twists, then $j(j-1)$ times elementary composition (1) for the braidings between V_i- and V_j-strands, and finally ij times elementary composition (1) for the braidings between V_i- and U_j-strands.

During the computation, we maintain a vector of size $(1 \times O(N^{i+j}))$. Applying Lemma 5.1, the sliding operation runs in $O(j(i+j)N^{i+j+2})$ arithmetic operations in $\mathcal{R}$, storing $O(N^{i+j})$ scalar from $\mathcal{R}$. In the algorithm, $i + j \leq \mathrm{cw}$, the carving-width of the link diagram. Consequently, we get $O(\mathrm{cw}^2\, N^{\mathrm{cw}+2})$ operations, with memory $O(N^{\mathrm{cw}})$.
Construction of evaluations and co-evaluations. The morphism $d_{U_1 \otimes \ldots \otimes U_k}$ appearing in Figure 7 is the result of k elementary compositions of type (4). The morphisms maintained during the computation are of size $(1 \times O(N^{2k}))$. Applying Lemma 5.1, the computation takes a total of $O(kN^{2k})$ arithmetic operations in $\mathcal{R}$, storing $O(N^{2k})$ scalars from $\mathcal{R}$. The case $b_{W_1 \otimes \ldots \otimes W_j}$ is similar.

In the algorithm, k (or j) is smaller than $\mathrm{cw}/2$. Consequently, the complexity is $O(\mathrm{cw}\, N^{\mathrm{cw}})$ arithmetic operations, storing $O(N^{\mathrm{cw}})$ scalars.

Composition of morphisms. Finally, the compositions of morphisms described in Figure 7 are implemented with a constant number of elementary compositions (5), (6), and (7). Considering Lemma 5.2, the product abc of dimensions never exceed $N^{\frac{3}{2}\,\mathrm{cw}}$. Consequently, the compositions of Figure 7 are implemented using $O(N^{\frac{3}{2}\,\mathrm{cw}})$ arithmetic operations in $\mathcal{R}$, storing $O(N^{\mathrm{cw}})$ scalars from $\mathcal{R}$.

Overall complexity. In conclusion, we sum up the different steps of the algorithm and its implementation. Let D be a coloured link diagram with n crossings and carving width cw, where the dimension of each colouring module is at most N. The algorithm first computes an optimal tree embedding in $O(\mathrm{poly}(n))$ operations. The tree has size n and width cw. W.l.o.g., we assume the diagram has at least one crossing that is not a twist, and consequently $\mathrm{cw} \geq 4$, the maximal degree of the graph. Considering $\mathrm{cw} \in O(\sqrt{n})$ and $\mathrm{cw} + 2 \leq \frac{3}{2}\,\mathrm{cw}$, the quantum invariant associated to the colouring is computed in:

$$O(n^2 N^{\frac{3}{2}\,\mathrm{cw}}) \text{ arithmetic operations in } \mathcal{R},$$

storing: $O(n)$ words for the diagram, plus $O(N^{\mathrm{cw}})$ scalars from $\mathcal{R}$.

6. Arithmetic complexity and quantum invariants of links

Working with matrices with $\mathcal{R}$-coefficients, for a ring $\mathcal{R}$, allows the algorithm to be applied in great generality. For example, any complex simple Lie algebra $\mathfrak{g}$ produces quantum invariants of links, that can be expressed as a composition of morphisms between free $\mathcal{R}$-modules, and to which our algorithm can be applied. See [26, Chapter 6] for an explicit construction.

In this case, $\mathcal{R}$ is a polynomial ring, and degrees of polynomials as well as values of coefficients may blow-up during intermediate computation. Specifically, both arithmetic operations within $\mathcal{R}$ and bit size of $\mathcal{R}$-elements may become *exponential in n*.

In this section we describe a solution to control the arithmetic complexity in the case $\mathcal{R} = \mathbb{Z}[q]$, which is sufficient for all $J_L^{\mathfrak{g}}$ invariants. We also provide detailed complexity bounds for completeness.

6.1. Arithmetic complexity of polynomial invariants.

We give coarse, but general, bounds on the degrees and coefficients of a polynomial invariant produced by the algorithm introduced above, that are sufficient for the complexity analysis.

Proposition 6.1. *Let $\mathcal{C}$ be a strict ribbon category of $\mathbb{Z}[q]$-modules, and let $D(L)$ be an n-crossings diagram of a link L whose components are coloured with free modules $V_1, \ldots, V_m \in \mathcal{C}$, of dimension at most N.*

Let d_0 and C_0 be respectively a bound on the degree and a bound on the absolute value of coefficients of all polynomials in the matrices $c^{\pm}_{V_i,V_j}$, $\theta^{\pm}_{V_i}$, d_{U_i}, b_{U_i}, for $1 \leq i, j \leq m$.

Then the polynomial invariant $J_L^{\mathcal{C}}(V_1, \ldots, V_m) \in \mathbb{Z}[q]$ has degree and absolute value of coefficients bounded by d_n and C_n respectively, with:

$$d_n = O(nd_0) \quad \textit{and} \quad C_n = 2^{O(n\sqrt{n}\log N + n\log C_0)}.$$

Proof. Consider a tree embedding of graph $D(L)$ where the tree is a path, with leaves attached to it. The minimal congestion over all such embeddings is called the *cut-width* of the graph, and is $O(\sqrt{n})$ due to the planar separator theorem.

Let k be the cut-width of $D(L)$, and $(\mathcal{P}, \phi)$ a minimal embedding of $D(L)$ into a path-tree. Running the algorithm of Sections 3-5 on this path decomposition boils down to computing the product of $O(n)$ matrices $M_{\alpha \cdot n} \cdot \ldots \cdot M_1$, where all matrices are tensor products of a

$c^{\pm}_{V_i,V_j}$, $\theta^{\pm}_{V_i}$, d_{U_i}, b_{U_i}, for some $1 \leq i, j \leq m$, with identities, and all matrices have size at most $N^{O(k)} \times N^{O(k)}$. Additionally, M_1 has 1 column, and $M_{\alpha \cdot n}$ has 1 row, to give a scalar in $\mathbb{Z}[q]$.

Tensor with the identity does not change the bounds d_0 and C_0 on degrees and coefficients. Multiplying by a matrix adds at most d_0 to the degree, and multiplies by at most $N^{O(k)}C_0$ the largest coefficient. We get the global bounds by multiplying the matrices together, and substituting $O(\sqrt{n})$ for k. □

We give a general algorithm to compute a one-variable, integer coefficient, polynomial invariants, using standard computer algebra techniques and the algorithm of Sections 3-5.

Proposition 6.2. *Let $\mathcal{C}$ be a strict ribbon category of $\mathbb{Z}[q]$-modules for the one-variable polynomial ring $\mathbb{Z}[q]$. Let L be an m-components link with colours the free modules $V_1, \ldots, V_m$, and let $J^{\mathcal{C}}_L(V_1, \ldots, V_m) \in \mathbb{Z}[q]$ be the associated topological invariant. Assume L is presented by an n-crossings diagram $D(L)$ with carving-width* cw.

Assume that the dimensions of the free modules $V_1, \ldots, V_m$ are at most N, and that the polynomial $J^{\mathcal{C}}_L(V_1, \ldots, V_m)$ has degree bounded by d_n and largest coefficient in absolute value bounded by C_n. Then $J^{\mathcal{C}}_L(V_1, \ldots, V_m)$ can be computed in:

$$O\Big(d_n\left(d_n + \log C_n\right) \cdot \mathrm{Ar}\left(\log(d_n \log d_n + \log C_n)\right) \times n N^{\frac{3}{2}\,\mathrm{cw}} \\ + d_n\left(d_n \log d_n + \log C_n\right)^2 + d_n^2\,\mathrm{Ar}\left(d_n \log d_n + \log C_n\right)\Big)$$

machine operations, using:

$$O\left(\log\left(d_n \log d_n + \log C_n\right) N^{\mathrm{cw}} + n d_n (d_n \log d_n + \log C_n) + d_n^2\,\mathrm{Ar}(d_n \log d_n + \log C_n)\right)$$

bits. Here, $\mathrm{Ar}(l) \in \widetilde{O}(l)$ is the arithmetic complexity of operations $+, -, \times, \div$ on integers encoded on at most l bits, which is linear in l up to a poly-logarithmic factor.

Proof. The algorithm relies on evaluation and interpolation. For short, denote $J^{\mathcal{C}}_L(V_1, \ldots, V_m)$ by $P(q) \in \mathbb{Z}[q]$.

Evaluation. We evaluate $P(q)$ on integer points $q \in \{0, 1, \ldots, d_n\}$. Fix q_0 in this set, and substitute q_0 for q in matrices $c^{\pm}_{V_i,V_j}$, $\theta^{\pm}_{V_i}$, d_{V_i}, and b_{V_i}. The algorithm of Sections 3-5 is consequently a succession of matrix multiplications, where all matrices have integer coefficients, and the resulting $P(q_0)$ is an integer of absolute value less than:

$$C d_n^{d_n+1} \leq 2^{(d_n+1)\log_2 d_n + \log_2 C_n} = 2^{O(d_n \log d_n + \log C_n)}$$

For a fixed q_0, we perform computation modulo the first r prime numbers $2 = p_1, \ldots, p_r$ successively, such that the product $p_1 \cdots p_r$ is larger than $|P(q_0)|$. We then reconstruct $P(q_0)$ using the Chinese Remainder Theorem. The product $p_1 \cdots p_r$ is of order $2^{r \log r}$ [23]. We take an appropriate r such that $r \log r \in \Theta(d_n \log d_n + \log C_n)$, which gives $r \in O(d_n + \log C_n)$.

Reconstructing the value $P(q_0)$ from all the $(P(q_0) \mod p_i)$, $1 \leq i \leq r$, can be computed in $O(r^2 \log^2 r) = O((d_n \log d_n + \log C_n)^2)$ machine operations [9, Theorem 5.8].

Additionally, the values of all primes p_i, $i \leq r$, are in $O(r \log(r \log r)) = O(r \log r) = O(d_n \log d_n + \log C_n)$ [23].

Denote by $\mathrm{Ar}(l)$ the computational complexity of performing arithmetic operations $+, -, \times$ on integers encoded on at most l bits, in $\mathbb{Z}/w\mathbb{Z}$, for an integer $w \leq 2^l$. The best known estimate for $C(l)$ is:

$$C(l) = O(l \log^2(l)\ 2^{O(\log^* l)}) = \widetilde{O}(l),$$

where $\log^*$ denotes the iterated logarithm, and the $\widetilde{O}$-notation hides poly-log factors. This describes the complexity of performing the extended Euclidean algorithm [9] using Fürer's method [6].

Interpolation. We reconstruct polynomial $P(q) \in \mathbb{Z}[q]$ of degree bounded by d_n using Lagrange interpolation. Lagrange interpolation gives directly a formula for $P(q)$, computable in $O(d_n^2 \operatorname{Ar}(d_n \log d_n + \log C_n))$ machine operations [9, Theorem 5.1].

Summing up the complexity of evaluating polynomial $P(q)$ on the first $d_n + 1$ non-negative integers using the modulo reconstruction approach and running the algorithm of Sections 3-5, and the complexity of evaluating the interpolation formula, gives the complexity of the proposition. □

We conclude by proving the main Theorem:

Proof. [of main Theorem 1.1] Fixing the category $\mathcal{C}$ and the colours $V_1, \ldots, V_m$, of dimension at most N, makes N constant, as well as the quantities d_0 and C_0 bounding degrees and coefficients of polynomials in the matrix for braidings, twists, and (co)evaluations. It enforces $d_n = O(n)$ (the bound on degree of the output polynomial), and $C_n = 2^{O(n\sqrt{n})}$ (the bound on absolute value of coefficients of the output invariant) in the complexity analysis. Substituting values gives the result of Theorem 1.1. □

Remark 6.3. Note that quantum invariants are usually defined in the category of $\mathbb{Z}[q, q^{-1}]$-modules. Multiplying the braiding, twist, and (co)evaluation matrices by q^a for a large enough, and re-normalising the output, allows us to restrict the algorithm to the case of $\mathbb{Z}[q]$-modules.

Note that we get the following parameterized complexity result for the more general problem of quantum invariant computation, where the invariant is part of the input:

Theorem 6.4. *The problem:*

> General quantum invariant problem:
> **Input**: *$\mathcal{C}, V_1, \ldots, V_m$, presented by braiding, twist, evaluation and co-evaluation matrices, and m-components link L, presented by a diagram $D(L)$,*
> **Output**: *quantum invariant $J_L^{\mathcal{C}}(V_1, \ldots, V_m)$*

can be solved in $O(\operatorname{poly}(n, d_0, \log C_0) N^{\frac{3}{2}\,\mathrm{cw}})$ machine operations, where n and cw *are respectively the number of crossings and the carving-width of the diagram $D(L)$, and d_0 and C_0 are respectively the maximal degree and maximal absolute value of coefficients of any polynomial in the input matrices.*

In other words, when the polynomials in the matrices are encoded with their lists of coefficients, the input size is $\Omega(\operatorname{poly}(N, d_0, \log C_0) + n)$, and the general quantum invariant problem is in the parameterized complexity class `XP`.

References

[1] Daniel Bienstock, *On embedding graphs in trees*, J. Comb. Theory, Ser. B **49** (1990), no. 1, 103–136.

[2] Benjamin Burton, *The next 350 million knots*, `https://regina-normal.github.io/data.html`, 2018.

[3] Benjamin A. Burton, *The HOMFLY-PT polynomial is fixed-parameter tractable*, 34th International Symposium on Computational Geometry, SoCG 2018, June 11-14, 2018, Budapest, Hungary, 2018, pp. 18:1–18:14.

[4] Benjamin A. Burton, Clément Maria, and Jonathan Spreer, *Algorithms and complexity for Turaev-Viro invariants*, J. Appl. Comput. Topol. **2** (2018), no. 1-2, 33–53.

[5] V.G. Drinfeld, *Quantum groups*, Proceedings International Congress of Mathematicians, 1986, pp. 798–820.
[6] Martin Fürer, *Faster integer multiplication*, SIAM J. Comput. **39** (2009), no. 3, 979–1005.
[7] Stavros Garoufalidis, *On the characteristic and deformation varieties of a knot*, Geometry & Topology Monographs **7** (2004), 291–309.
[8] ———, *The Jones slopes of a knot*, Quantum Topology **2** (2011), no. 1, 43–69.
[9] Joachim Von Zur Gathen and Jurgen Gerhard, *Modern computer algebra*, 2 ed., Cambridge University Press, New York, NY, USA, 2003.
[10] Qian-Ping Gu and Hisao Tamaki, *Optimal branch-decomposition of planar graphs in o(n3) time*, ACM Trans. Algorithms **4** (2008), no. 3, 30:1–30:13.
[11] Kristóf Huszár and Jonathan Spreer, *3-manifold triangulations with small treewidth*, 35th International Symposium on Computational Geometry, SoCG 2019, June 18-21, 2019, Portland, Oregon, USA., 2019, pp. 44:1–44:20.
[12] Kristóf Huszár, Jonathan Spreer, and Uli Wagner, *On the treewidth of triangulated 3-manifolds*, 34th International Symposium on Computational Geometry, SoCG 2018, June 11-14, 2018, Budapest, Hungary, 2018, pp. 46:1–46:15.
[13] F. Jaeger, D. L. Vertigan, and D. J. A. Welsh, *On the computational complexity of the Jones and Tutte polynomials*, Mathematical Proceedings of the Cambridge Philosophical Society **108** (1990), no. 1, 3553.
[14] Michio Jimbo, *A q-difference analogue of U(g) and the Yang-Baxter equation*, Letters in Mathematical Physics **10** (1985), no. 1, 63–69.
[15] Vaughan F. R. Jones, *A polynomial invariant for knots via von Neumann algebras*, Bull. Amer. Math. Soc. **12** (1985), 103–111.
[16] R. M. Kashaev, *The hyperbolic volume of knots from the quantum dilogarithm*, Letters in Mathematical Physics **39** (1997), no. 3, 269–275.
[17] W. B. Raymond Lickorish, *An introduction to knot theory*, Graduate Texts in Mathematics, vol. 175, Springer-Verlag New York, 1997.
[18] J. A. Makowsky and J. P. Mariño, *The parameterized complexity of knot polynomials*, J. Comput. Syst. Sci. **67** (2003), no. 4, 742–756.
[19] Clément Maria and Jessica S. Purcell, *Treewidth, crushing, and hyperbolic volume*, Algebraic & Geometric Topology **19** (2019), 2625–2652.
[20] Hitoshi Murakami and Jun Murakami, *The colored Jones polynomials and the simplicial volume of a knot*, Acta Math. **186** (2001), no. 1, 85–104.
[21] N. Yu. Reshetikhin and V. G. Turaev, *Ribbon graphs and their invariants derived from quantum groups*, Comm. Math. Phys. **127** (1990), no. 1, 1–26.
[22] Neil Robertson and P. D. Seymour, *Graph minors. II. Algorithmic aspects of tree-width*, J. Algorithms **7** (1986), no. 3, 309–322.
[23] J. B. Rosser and L. Schoenfeld, *Approximate formulas for some functions of prime numbers*, IJM **6** (1962), 64–94. MR 25, 1139
[24] Saul Schleimer, Arnaud de Mesmay, Jessica Purcell, and Eric Sedgwick, *On the tree-width of knot diagrams*, Journal of Computational Geometry **10** (2019), no. 1, 164–180.
[25] P. D. Seymour and R. Thomas, *Call routing and the ratcatcher*, Combinatorica **14** (1994), no. 2, 217–241.
[26] Vladimir G. Turaev, *Quantum invariants of knots and 3-manifolds*, revised ed., de Gruyter Studies in Mathematics, vol. 18, Walter de Gruyter & Co., Berlin, 2010.
[27] Vladimir G. Turaev and Oleg Y. Viro, *State sum invariants of* 3*-manifolds and quantum* 6j*-symbols*, Topology **31** (1992), no. 4, 865–902.

INRIA Sophia Antipolis-Méditerranée, FRANCE
E-mail address: `clement.maria@inria.fr`