

A characterization of indecomposable web modules over Khovanov–Kuperberg algebras

LOUIS-HADRIEN ROBERT

After shortly reviewing the construction of the Khovanov–Kuperberg algebras, we give a characterization of indecomposable web modules. It says that a web module is indecomposable if and only if one can deduce its indecomposability directly from the Kuperberg bracket (via a Schur lemma argument). The proof relies on the construction of idempotents given by explicit foams. These foams are encoded by combinatorial data called red graphs. The key point is to show that when the Schur lemma does not apply for a web w , an appropriate red graph for w can be found.

17B37; 57M27, 57R56

1 Introduction

In this paper, we study the category of projective modules over Khovanov–Kuperberg algebras. These algebras emerge naturally when one extends algebraically the \mathfrak{sl}_3 –homology to tangles.

The \mathfrak{sl}_3 –homology, defined by Khovanov [5], is a categorification of the \mathfrak{sl}_3 –invariant, which is itself a specialization of the HOMFLY-PT polynomial; see Freyd et al [2] and Przytycki and Traczyk [13]. Mackaay and Vaz [11] and Morrison and Nieh [12] extended the \mathfrak{sl}_3 –homology to tangles in a Bar-Natan fashion. More recently, Lauda, Queffelec and Rose [9] considered such an extension, and understood it as arising from the 2–representation theory of the categorified quantum \mathfrak{sl}_m ’s.

For any¹ finite sequence of signs ε , the Khovanov–Kuperberg algebra K^ε is an \mathfrak{sl}_3 –analogue of Khovanov’s arc algebras H^n [4]. They were introduced in 2012 by Mackaay, Pan and Tubbenhauer [10] and the author [14].

Each algebra K^ε contains a family of orthogonal idempotents (1_w) (w runs through the set of isotopy classes of nonelliptic ε –webs; see Kuperberg [8]). In contrast with the \mathfrak{sl}_2 case, these idempotents do not need to be primitive. In [10], Mackaay, Pan

¹Actually, the sequences of signs need to be admissible. This means that the sum of signs is equal to 0 modulo 3. This restriction is the \mathfrak{sl}_3 –analogue of the parity restriction for Khovanov’s arcs algebras H^n .

and Tubbenhauer proved that the projective modules associated with these idempotents form a base of the Grothendieck group of $K^\varepsilon - \text{proj}_{\text{gr}}$. The author [14] provided a combinatorial condition for a module corresponding to a nonelliptic web to be indecomposable.

The aim of this paper is to provide a handy criterion for an idempotent 1_w to be primitive.

Theorem *Let w be a nonelliptic ε -web. The idempotent 1_w is indecomposable if and only if $\langle \bar{w}w \rangle$ is monic of degree $l(\varepsilon)$.*

In this statement, $l(\varepsilon)$ represents the length of ε and $\langle \bar{w}w \rangle$ represents the Kuperberg bracket [8] of the web obtained by gluing w and its reoriented mirror image along ε . The implication from right to left is immediate (see Lemma 2.17).

We would like to point out a few facts about the proof.

- (1) The proof is fully constructive: if an idempotent 1_w is not primitive, we give two explicit orthogonal idempotents e_1 and e_2 such that $1_w = e_1 + e_2$.
- (2) We develop a handy framework for certain calculations with foams (Section 3B).
- (3) If w is a nonelliptic ε -web, we prove that a (w, w) -foam is equivalent to a nonzero multiple of $I \times w$ if and only if it is isotopic to $I \times w$.

One can construct the algebras K^ε over any unital commutative ring; in this paper we work over \mathbb{Q} . If one extends the coefficients of the algebras K^ε to \mathbb{C} , the main theorem follows from the paper of Mackaay, Pan and Tubbenhauer [10], where they proved in this context that the base given by nonelliptic webs is unitriangular with respect to the dual canonical basis.

The proof of the main theorem relies mainly on one combinatorial tool called a *red graph*. These are subgraphs of the dual graph of a web satisfying certain conditions. The proof proceeds in two distinct steps. First, we prove that given a red graph, we can construct a nontrivial idempotent. Then we show that whenever $\langle \bar{w}w \rangle$ is not monic of degree $l(\varepsilon)$, we can find a red graph for w .

Organization of the paper

In Section 2, we explain the construction of the Khovanov–Kuperberg algebras. In Section 3, we introduce red graphs and construct idempotents from red graphs. We prove in Section 3C that the constructed idempotents are not trivial. Section 4 is the longest and the most technical. The main point is to give an algorithm to construct a red graph in a web such that $\langle \bar{w}w \rangle$ is not monic of degree $l(\varepsilon)$. This requires the introduction of stacks of red graphs and a way to recover red graphs from these stacks.




Figure 1: Example of a closed web

Acknowledgments The author is indebted to the referee for pointing out numerous typos in a previous version of this paper and suggesting many improvements. He wishes to thank Christian Blanchet for suggesting the subject.

2 The Khovanov–Kuperberg algebras

2A The 2–category of (\cdot, \cdot) –webs

2A.1 Webs In the following $\varepsilon = (\varepsilon^1, \dots, \varepsilon^n)$ (or $\varepsilon_0, \varepsilon_1$ etc) will always be a finite sequence of signs, its length n will be denoted by $l(\varepsilon)$, and such an ε will be *admissible* if $\sum_{i=1}^{l(\varepsilon)} \varepsilon^i$ is divisible by 3.

Definition 2.1 (Kuperberg [8]) A *closed web* is a plane trivalent oriented finite graph (with possibly some vertexless loops and multiple edges) such that every vertex is either a sink or a source.

Remark The orientation condition is equivalent to saying that the graph is bipartite (by sinks and sources). Given any oriented embedded disk in \mathbb{R}^2 which transversely intersects a web, the algebraic intersection number of the web with this disk is equal to 0 modulo 3. To express this fact, we say that the *flow* of a web is preserved modulo 3.

Proposition 2.2 A closed web contains at least a square, a digon or a vertexless circle.

Proof It is enough to consider a connected web w . A connected web is always 2–connected (because of the flow); hence it makes sense to use the Euler characteristic. Suppose that w is not a circle. Then

$$\#F - \#E + \#V = 2,$$

but $3\#V = 2\#E$. And if we denote by F_i the set of faces with i sides, we have

$$\sum_{i>0} i\#F_i = 2\#E.$$

Altogether, this gives

$$\sum_{i>0} \#F_i - \frac{i}{6} \#F_i = 2,$$

which proves that some faces have strictly less than 6 sides. □

Proposition 2.3 (Kuperberg [8]) *There exists one and only one map $\langle \cdot \rangle$ from closed webs to Laurent polynomials in q which is invariant by isotopy, multiplicative with respect to disjoint union and which satisfies the local relations*

$$\begin{aligned} \langle \text{square with arrows} \rangle &= \langle \text{cup} \rangle \langle \text{cap} \rangle + \langle \text{cup} \rangle \langle \text{cap} \rangle, \\ \langle \text{diamond} \rangle &= [2] \cdot \langle \text{cup} \rangle, \\ \langle \text{circle} \rangle &= \langle \text{circle} \rangle = [3], \end{aligned}$$

where $[n] \stackrel{\text{def}}{=} (q^n - q^{-n}) / (q - q^{-1})$. We call this polynomial the Kuperberg bracket. The formulas being symmetric in q and q^{-1} , the Kuperberg bracket is a symmetric Laurent polynomial.

Proof Uniqueness comes from the remark on page 1305; existence follows from the representation-theoretic point of view developed by Kuperberg [8] and means the given relations are consistent. A nonquantified version of this result is due to Jaeger [3]. □

Definition 2.4 The *degree* of a symmetric Laurent polynomial $P(q) = \sum_{i \in \mathbb{Z}} a_i q^i$ is $\max_{i \in \mathbb{Z}} \{i \text{ such that } a_i \neq 0\}$.

Definition 2.5 A $(\varepsilon_0, \varepsilon_1)$ -web w is an intersection of a closed web w' with $[0, 1] \times [0, 1]$ such that:

- There exists $\eta_0 \in]0, 1]$ such that

$$w \cap [0, 1] \times [0, \eta_0] = \left\{ \frac{1}{2l(\varepsilon_0)}, \frac{1}{2l(\varepsilon_0)} + \frac{1}{l(\varepsilon_0)}, \frac{1}{2l(\varepsilon_0)} + \frac{2}{l(\varepsilon_0)}, \dots, \frac{1}{2l(\varepsilon_0)} + \frac{l(\varepsilon_0) - 1}{l(\varepsilon_0)} \right\} \times [0, \eta_0].$$

- There exists $\eta_1 \in [0, 1[$ such that

$$w \cap [0, 1] \times [\eta_1, 1] = \left\{ \frac{1}{2l(\varepsilon_1)}, \frac{1}{2l(\varepsilon_1)} + \frac{1}{l(\varepsilon_1)}, \frac{1}{2l(\varepsilon_1)} + \frac{2}{l(\varepsilon_1)}, \dots, \frac{1}{2l(\varepsilon_1)} + \frac{l(\varepsilon_1) - 1}{l(\varepsilon_1)} \right\} \times [\eta_1, 1].$$

- The orientations of the edges of w match $-\varepsilon_0$ and $+\varepsilon_1$ (see Figure 2 for conventions).




Figure 2: Two examples of $(\varepsilon_0, \varepsilon_1)$ -webs with $\varepsilon_0 = (+, +, +)$ and $\varepsilon_1 = (+, +, -, -)$ and the mirror image of the second one

When ε_1 is the empty sequence we speak of ε_0 -webs, we call ε the *boundary* of an ε -web, and when we do not want to specify ε_1 and ε_2 we speak of (\cdot, \cdot) -webs.

If w_1 is an $(\varepsilon_0, \varepsilon_1)$ -web and w_2 is an $(\varepsilon_1, \varepsilon_2)$ -web, we define $w_1 w_2$ to be the $(\varepsilon_0, \varepsilon_2)$ -web obtained by gluing w_1 and w_2 along ε_1 and resizing. Note that this can be thought of as a composition if we think about an $(\varepsilon, \varepsilon')$ -web as a morphism from ε' to ε (ie a (\cdot, \cdot) -web should be read as a morphism from top to bottom). The *mirror image* of an $(\varepsilon_0, \varepsilon_1)$ -web w is the mirror image of w with respect to $\mathbb{R} \times \{\frac{1}{2}\}$ with all orientations reversed. This is an $(\varepsilon_1, \varepsilon_0)$ -web and we denote it by \bar{w} . If w is an $(\varepsilon, \varepsilon)$ -web the *closure* of w is the closed web obtained by connecting the top and the bottom by simple arcs (this is like a braid closure). We denote it by $\text{tr}(w)$.

Definition 2.6 A (\cdot, \cdot) -web with no circle, no digon and no square is said to be *nonelliptic*. The nonelliptic (\cdot, \cdot) -webs are the minimal ones in the sense that they cannot be reduced by the relations of Proposition 2.3.

Proposition 2.7 (Kuperberg [8]) *For any given couple $(\varepsilon_0, \varepsilon_1)$ of sequences of signs there are finitely many nonelliptic $(\varepsilon_0, \varepsilon_1)$ -webs.*

Remark From the combinatorial flow modulo 3, we obtain that there exist some $(\varepsilon_0, \varepsilon_1)$ -webs if and only if the sequence $-\varepsilon_0$ concatenated with ε_1 is admissible.

2A.2 Foams All material here comes from the paper of Khovanov [5].

Definition 2.8 A *prefoam* is a smooth oriented compact surface Σ (its connected component are called *facets*) together with the following data:

- A partition of the connected components of the boundary into cyclically ordered 3-sets and for each 3-set (C_1, C_2, C_3) , three orientation preserving diffeomorphisms $\phi_1: C_2 \rightarrow C_3$, $\phi_2: C_3 \rightarrow C_1$ and $\phi_3: C_1 \rightarrow C_2$ such that $\phi_3 \circ \phi_2 \circ \phi_1 = \text{id}_{C_2}$.
- A function from the set of facets to \mathbb{N} (this gives the number of *dots* on each facet).




Figure 3: Singularities of a prefoam

The CW complex associated with a prefoam is the 2–dimensional CW complex Σ quotiented by the diffeomorphisms so that the three circles of one 3–set are identified and become just one, called a *singular circle*. The *degree* of a prefoam f is equal to $-2\chi(\Sigma')$, where χ is the Euler characteristic and Σ' is the CW complex associated with f with the dots punctured out (ie a dot increases the degree by 2).

Remark A CW complex associated with a prefoam has two local models depending on whether we are on a singular circle or not. If a point x is not on a singular circle, then it has a neighborhood diffeomorphic to a 2–dimensional disk, else it has a neighborhood diffeomorphic to a Y shape times an interval (see Figure 3).

Definition 2.9 A *closed foam* is the image of an embedding of the CW complex associated with a prefoam such that the cyclic orders of the prefoam are compatible with the left hand rule in \mathbb{R}^3 with respect to the orientations of the singular circles.² The *degree* of a closed foam is the degree of the underlying prefoam.

Definition 2.10 If w_b and w_t are $(\varepsilon_0, \varepsilon_1)$ –webs, a (w_b, w_t) –foam f is the intersection of a foam f' with $\mathbb{R} \times [0, 1] \times [0, 1]$ such that:

- There exists $\eta_0 \in]0, 1]$ such that

$$f \cap \mathbb{R} \times [0, \eta_0] \times [0, 1] = \left\{ \frac{1}{2l(\varepsilon_0)}, \frac{1}{2l(\varepsilon_0)} + \frac{1}{l(\varepsilon_0)}, \frac{1}{2l(\varepsilon_0)} + \frac{2}{l(\varepsilon_0)}, \dots, \frac{1}{2l(\varepsilon_0)} + \frac{l(\varepsilon_0) - 1}{l(\varepsilon_0)} \right\} \times [0, \eta_0] \times [0, 1].$$

- There exists $\eta_1 \in [0, 1[$ such that

$$f \cap \mathbb{R} \times [\eta_1, 1] \times [0, 1] = \left\{ \frac{1}{2l(\varepsilon_1)}, \frac{1}{2l(\varepsilon_1)} + \frac{1}{l(\varepsilon_1)}, \frac{1}{2l(\varepsilon_1)} + \frac{2}{l(\varepsilon_1)}, \dots, \frac{1}{2l(\varepsilon_1)} + \frac{l(\varepsilon_1) - 1}{l(\varepsilon_1)} \right\} \times [\eta_1, 1] \times [0, 1].$$

²We mean here that if, next to a singular circle, with the forefinger of the left hand we go from face 1 to face 2 to face 3 the thumb points to indicate the orientation of the singular circle (induced by orientations of facets). This is not quite canonical, physicists prefer the right-hand rule, but this is the convention used by Khovanov [5].

- There exists $\eta_b \in]0, 1]$ such that $f \cap \mathbb{R} \times [0, 1] \times [0, \eta_b] = w_b \times [0, \eta_b]$.
- There exists $\eta_t \in [0, 1[$ such that $f \cap \mathbb{R} \times [0, 1] \times [\eta_t, 1] = w_t \times [\eta_t, 1]$, where the orientations of the facets of f are compatible with the those of w_t and the reversed orientations of w_b .

The *degree* of a (w_b, w_t) -foam f is equal to $\chi(w_b) + \chi(w_t) - 2\chi(\Sigma)$, where Σ is the underlying CW complex associated with f with the dots punctured out.

If f_b is a (w_b, w_m) -foam and f_t is a (w_m, w_t) -foam we define $f_b f_t$ to be the (w_b, w_t) -foam obtained by gluing f_b and f_t along w_m and resizing. This operation may be thought of as a composition if we think of a (w_1, w_2) -foam as a morphism from w_2 to w_1 ; ie from the top to the bottom. This composition map is degree preserving. Like for the webs, we define the *mirror image* of a (w_1, w_2) -foam f to be the (w_2, w_1) -foam which is the mirror image of f with respect to $\mathbb{R} \times \mathbb{R} \times \{\frac{1}{2}\}$ with all orientations reversed. We denote it by \bar{f} .

Definition 2.11 If $\varepsilon_0 = \varepsilon_1 = \emptyset$ and w is a closed web, then a (\emptyset, w) -foam is simply called *foam* or w -foam when one wants to focus on the boundary of the foam.

All these data together lead to the definition of a monoidal 2-category.

Definition 2.12 The 2-category \mathcal{WT} is the monoidal³ 2-category where

- the objects are finite sequences of signs,
- the 1-morphisms from ε_1 to ε_0 are isotopy classes (with fixed boundary) of $(\varepsilon_0, \varepsilon_1)$ -webs,
- the 2-morphisms from \widehat{w}_t to \widehat{w}_b are \mathbb{Q} -linear combinations of isotopy classes of (w_b, w_t) -foams, where $\widehat{\cdot}$ stands for “the isotopy class of”. The 2-morphisms come with a grading and the composition respects the degree.

The monoidal structure is given by concatenation of sequences at the 0 level, and disjoint union of vertical strands or disks (with corners) at the 1 and 2 levels.

Remark Note that the 2-morphism spaces contain morphisms of arbitrary degree. To emphasize this specifically, we denote the hom-space between two (\cdot, \cdot) -webs by HOM rather than hom which usually refers to spaces of 2-morphisms of degree 0.

³Here we choose a rather strict point of view and hence the monoidal structure is strict (we consider everything up to isotopy), but it is possible to define the notion in a nonstrict context, and the same data leads to a monoidal bicategory.

2B Khovanov’s TQFT for (\cdot, \cdot) -webs

In [5], Khovanov defines a numerical invariant for prefoams and this allows him to construct a TQFT \mathcal{F} from the category $\text{HOM}_{\mathcal{WT}}(\emptyset, \emptyset)$ to the category of graded \mathbb{Q} -modules, (via a universal construction à la Blanchet–Habegger–Masbaum–Vogel [1]). This TQFT is graded (this comes from the fact that prefoams with nonzero degree are evaluated to zero), and satisfies the following local relations (brackets indicate grading shifts):

$$\begin{aligned} \mathcal{F}\left(\begin{array}{c} \text{---} \text{---} \text{---} \\ \diagup \quad \diagdown \\ \text{---} \text{---} \text{---} \\ \diagdown \quad \diagup \\ \text{---} \text{---} \text{---} \end{array}\right) &= \mathcal{F}\left(\begin{array}{c} \text{---} \text{---} \\ \diagup \quad \diagdown \\ \text{---} \end{array}\right) \oplus \mathcal{F}\left(\begin{array}{c} \text{---} \text{---} \\ \diagdown \quad \diagup \\ \text{---} \end{array}\right), \\ \mathcal{F}\left(\begin{array}{c} \text{---} \\ \diagup \quad \diagdown \\ \text{---} \\ \diagdown \quad \diagup \\ \text{---} \end{array}\right) &= \mathcal{F}\left(\begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \end{array}\right)\{-1\} \oplus \mathcal{F}\left(\begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \end{array}\right)\{1\}, \\ \mathcal{F}\left(\begin{array}{c} \text{---} \\ \text{---} \end{array}\right) &= \mathcal{F}\left(\begin{array}{c} \text{---} \\ \text{---} \end{array}\right) = \mathbb{Q}\{-2\} \oplus \mathbb{Q} \oplus \mathbb{Q}\{2\}. \end{aligned}$$

These relations show that \mathcal{F} is a categorified counterpart of the Kuperberg bracket. We sketch the construction below.

Definition 2.13 We denote by \mathcal{A} the Frobenius algebra $\mathbb{Z}[X]/(X^3)$ with trace τ given by

$$\tau(X^2) = -1, \quad \tau(X) = 0, \quad \tau(1) = 0.$$

We equip \mathcal{A} with a grading by setting $\text{deg}(1) = -2$, $\text{deg}(X) = 0$ and $\text{deg}(X^2) = 2$. With these settings, the multiplication has degree 2 and the trace has degree -2 . The co-multiplication is determined by the multiplication and the trace and we have

$$\begin{aligned} \Delta(1) &= -1 \otimes X^2 - X \otimes X - X^2 \otimes 1, \\ \Delta(X) &= -X \otimes X^2 - X^2 \otimes X, \\ \Delta(X^2) &= -X^2 \otimes X^2. \end{aligned}$$

This Frobenius algebra gives us a $1 + 1$ TQFT (this is well known, see for example the book of Kock [7] for details), we denote it by \mathcal{F} : the circle is sent to \mathcal{A} , a cup to the unity, a cap to the trace, and a pair of pants either to multiplication or co-multiplication. A dot on a surface will represent multiplication by X so that \mathcal{F} extends to the category of oriented dotted $(1 + 1)$ -cobordisms. We have a surgery formula given by Figure 4.

This TQFT gives of course a numerical invariant for closed dotted oriented surfaces. If we define numerical values for the differently dotted theta prefoams (the theta prefoam consists of three disks with trivial diffeomorphisms between their boundaries; see




Figure 4: The surgery formula for the TQFT \mathcal{F}




Figure 5: The dotless theta prefoam



Figure 6: The evaluations of dotted theta foams; the evaluation is unchanged when one cyclically permutes the faces. All the configurations which cannot be obtained from these by cyclic permutation are sent to 0 by \mathcal{F} .

Figure 5), then by applying the surgery formula, we can compute a numerical value for prefoams.

In [5], Khovanov shows that setting the evaluations of the dotted theta foams, as shown in Figure 6, leads to a well-defined numerical invariant \mathcal{F} for prefoams. This numerical invariant gives the opportunity to build a (closed web, (\cdot, \cdot) –foams)–TQFT: for a given web w , consider the graded \mathbb{Q} –module generated by all the (w, \emptyset) –foams, and mod this space out by the kernel of the bilinear map $(f, g) \mapsto \mathcal{F}(\bar{f}g)$. Note that $\bar{f}g$ is a closed foam. Khovanov showed that the obtained graded vector spaces are finite-dimensional with graded dimensions given by the Kuperberg formulas, and he showed that we have the local relations described in Figure 7.

This method allows us to define a new graded 2–category $\mathbb{W}\mathbb{T}$. Its objects and its 1–morphisms are those of the 2–category $\mathcal{W}\mathcal{T}$ while its 2–morphism spaces are those of $\mathcal{W}\mathcal{T}$ modded out like in the last paragraph. One should notice that a (w_b, w_t) –foam can always be deformed into a $(\text{tr}(\bar{w}_b w_t), \emptyset)$ –foam and vice-versa. Khovanov’s results restated in this language give that if w_b and w_t are $(\varepsilon_0, \varepsilon_1)$ –webs, the graded dimension of $\text{HOM}_{\mathbb{W}\mathbb{T}}(w_t, w_b)$ is given by

$$(1) \quad \langle \text{tr}(\bar{w}_b w_t) \rangle \cdot q^{l(\varepsilon_0) + l(\varepsilon_1)}.$$

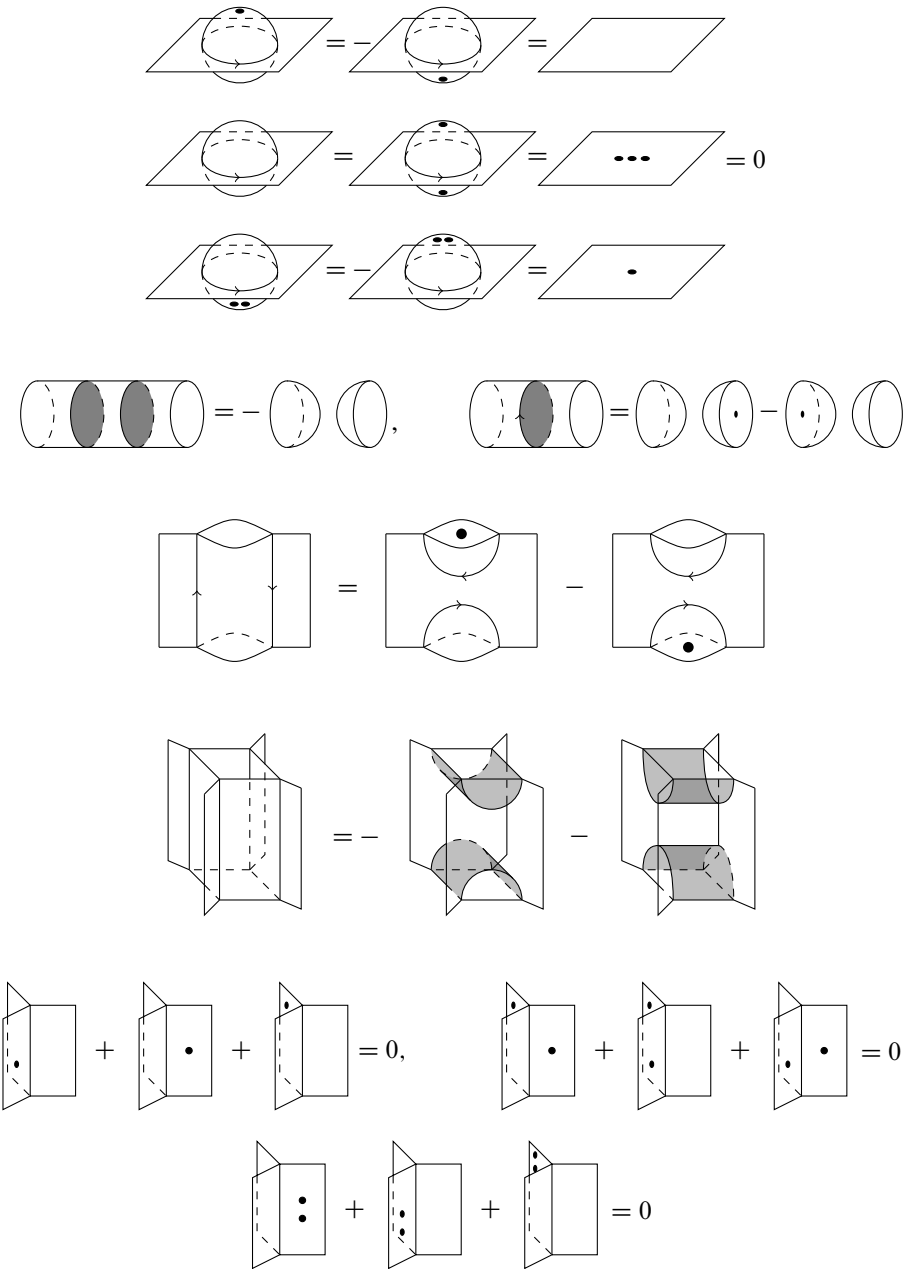


Figure 7: Local relations for 2-morphisms in WT. The first 3 lines are called bubbles relations, the next is the bamboo relations, the one after is the digon relation, then we have the square relation and finally we show the 3 dot migration relations.

Note that when $\varepsilon_1 = \emptyset$, there is no need to take the closure, because $w_b \bar{w}_t$ is already a closed web. The shift by $l(\varepsilon_0) + l(\varepsilon_1)$ comes from the fact that $\chi(\text{tr}(\bar{w}_b w_t)) = \chi(w_t) + \chi(w_b) - (l(\varepsilon_0) + l(\varepsilon_1))$.

Definition 2.14 We consider the set of local relations which consists of

- the surgery relation,
- the evaluations of the dotted spheres and of the dotted theta foams,
- the square relations and the digon relations (see [Figure 7](#)).

We call them the *foam relations*.

2C The Khovanov–Kuperberg algebra K^ε

We want to extend the Khovanov TQFT to the 0–dimensional objects, ie to build a 2–functor from the 2–category \mathcal{WT} to the 2–category of algebras. We follow the methodology of Khovanov [4] and we start by defining the image of the 0–objects: the algebras K^ε . This can be compared with the paper of Mackaay, Pan and Tubbenhauer [10].

Definition 2.15 Let ε be an admissible finite sequence of signs. We define \tilde{K}^ε to be the full subcategory of $\text{HOM}_{\mathbb{WT}}(\emptyset, \varepsilon)$ whose objects are nonelliptic ε –webs. This is a graded \mathbb{Q} –algebroid. We recall that a k –algebroid is a k –linear category. This can be seen as an algebra by setting

$$K^\varepsilon = \bigoplus_{(w_b, w_t) \in (\text{ob}(\tilde{K}^\varepsilon))^2} \text{HOM}_{\mathbb{WT}}(w_b, w_t)$$

and the multiplication on K^ε is given by the composition of morphisms in \tilde{K}^ε whenever it’s possible and by zero when it’s not. We write

$${}_{w_1} K^\varepsilon {}_{w_2} \stackrel{\text{def}}{=} \text{HOM}_{\mathbb{WT}}(w_2, w_1).$$

This is a unitary algebra because of [Proposition 2.7](#). The unit element is $\sum_{w \in \text{ob}(\tilde{K}^\varepsilon)} 1_w$. Suppose ε is fixed; for a nonelliptic ε –web w we define P_w to be the left K^ε –module

$$P_w = \bigoplus_{w' \in \text{ob}(\tilde{K}^\varepsilon)} \text{HOM}_{\mathbb{WT}}(w, w') = \bigoplus_{w' \in \text{ob}(\tilde{K}^\varepsilon)} {}_{w'} K^\varepsilon {}_w.$$

The structure of module is given by composition on the left.

For a given ε , the modules P_w are all projective and we have the following decomposition in the category of left K^ε -modules:

$$K^\varepsilon = \bigoplus_{w \in \text{ob}(\tilde{K}_\varepsilon)} P_w.$$

Proposition 2.16 *Let ε be an admissible sequence of signs, and w_1 and w_2 two nonelliptic ε -webs. Then the graded dimension of $\text{HOM}_{K^\varepsilon}(P_{w_1}, P_{w_2})$ is given by $\langle (\bar{w}_1 w_2) \rangle \cdot q^{l(\varepsilon)}$.*

Proof An element of $\text{HOM}_{K^\varepsilon}(P_{w_1}, P_{w_2})$ is completely determined by the image of 1_{w_1} and this element can be sent to any element of $\text{HOM}_{\mathbb{W}\mathbb{T}}(P_{w_2}, P_{w_1})$, and $\dim_q(\text{HOM}_{\mathbb{W}\mathbb{T}}(P_{w_1}, P_{w_2})) = \langle (\bar{w}_1 w_2) \rangle \cdot q^{l(\varepsilon)}$ thanks to (1). \square

In what follows we will prove that certain modules are indecomposable. They all have finite dimension over \mathbb{Q} hence it's enough to show that their rings of endomorphisms contain no nontrivial idempotents. Idempotents always have degree zero, so we have the following lemma.

Lemma 2.17 *If w is a nonelliptic ε -web such that $\langle \bar{w} w \rangle$ is monic of degree $l(\varepsilon)$, then the graded K^ε -module P_w is indecomposable.*

Proof This follows from the previous discussion: the hypothesis on $\langle \bar{w} w \rangle$ implies that $\dim((\text{HOM}_{\mathbb{W}\mathbb{T}}(P_w, P_w)_0) = 1$. This shows that $\text{HOM}_{\mathbb{W}\mathbb{T}}(P_w, P_w)_0$ contains only multiples of the identity and hence that it contains no nontrivial idempotent. \square

A similar lemma proves that two modules are not isomorphic.

Lemma 2.18 *If w_1 and w_2 are two nonelliptic ε -webs such that $\langle \bar{w}_1 w_2 \rangle$ has degree strictly smaller than $l(\varepsilon)$, then the graded K^ε -modules P_{w_1} and P_{w_2} are not isomorphic.*

Remark In [10], Mackaay, Pan and Tubbenhauer proved that if w_1 and w_2 are two nonisotopic webs, then P_{w_1} and P_{w_2} are not isomorphic, which is a stronger result. In contrast with the main theorem, their result remains true when working over \mathbb{Q} instead of over \mathbb{C} .

Proof If they were isomorphic, there would exist two morphisms f and g such that $f \circ g = 1_{P_{w_1}}$ and therefore $f \circ g$ would have degree zero. The hypothesis of the lemma implies that f and g (because $\langle \bar{w}_1 w_2 \rangle = \langle \bar{w}_2 w_1 \rangle$) have positive degree, so the degree of their composite is positive too. \square

Remark The way we constructed the algebra K^ε is very similar to the construction of H^n in [4] by Khovanov. Using the same method we can finish the construction of a $0 + 1 + 1$ TQFT:

- If ε is an admissible sequence of signs, then $\mathcal{F}(\varepsilon) = K^\varepsilon$.
- If w is an $(\varepsilon_1, \varepsilon_2)$ -web with ε_1 and ε_2 admissible, then

$$\mathcal{F}(w) = \bigoplus_{\substack{u \in \text{ob}(\tilde{K}^{\varepsilon_1}) \\ v \in \text{ob}(\tilde{K}^{\varepsilon_2})}} \mathcal{F}(\bar{u}wv),$$

and it has a structure of graded K^{ε_1} -module- K^{ε_2} (ie $(K^{\varepsilon_1}, K^{\varepsilon_2})$ -bimodule). Note that if w is a nonelliptic ε -web, then $\mathcal{F}(w) = P_w$.

- If w and w' are two $(\varepsilon_1, \varepsilon_2)$ -webs and f is a (w, w') -foam, then we set

$$\mathcal{F}(f) = \sum_{\substack{u \in \text{ob}(\tilde{K}^{\varepsilon_1}) \\ v \in \text{ob}(\tilde{K}^{\varepsilon_2})}} \mathcal{F}(\bar{u}f_v),$$

where $\bar{u}f_v$ is the foam f with glued on its sides $\bar{u} \times [0, 1]$ and $v \times [0, 1]$. This is a map of graded K^{ε_1} -modules- K^{ε_2} .

We encourage the reader to have a look at the beautiful construction for the \mathfrak{sl}_2 case in [4].

In the \mathfrak{sl}_2 case the classification of projective indecomposable modules is fairly easy, because all the all web modules associated with arc diagrams without circles are indecomposable. In our case, this situation is much more complicated because of the following proposition.

Proposition 2.19 ([6; 12], see [15] for details) *Let ε be the sequence of signs $(+, -, -, +, +, -, -, +, +, -, -, +)$ and let w and w_0 be the two ε -webs given by Figure 8. Then the web module P_w is decomposable and admits P_{w_0} as a direct factor.*

The aim of this paper is to prove the following statement which is a converse to Lemma 2.17.

Theorem 2.20 *Let w be an ε -web. The K^ε -module P_w is indecomposable if and only if $\langle \bar{w}w \rangle$ is monic of degree $l(\varepsilon)$. Furthermore if the K^ε -module P_w is decomposable, it contains another web module as a direct factor.*




Figure 8: The ε -webs w (left) and w_0 (right). To fit into the formal context of the 2-category one should stretch the outside edges to horizontal line below the whole picture. We draw it this way to enjoy more symmetry. To simplify we didn't draw the arrows.

The proof relies on a combinatorial tool called red graphs. In Section 3, we give an explicit construction (in terms of foams) of a nontrivial idempotent associated to a red graph. In Section 4, we show that when an ε -web w is such that $\langle \bar{w}w \rangle$ is not monic of degree $l(\varepsilon)$, it contains a red graph.

3 Idempotent from red graphs

3A Definitions

Definition 3.1 Let w be an ε -web. A red graph for w is a nonempty subgraph G of the dual graph of w denoted by $D(w)$ such that:

- (i) All faces belonging to $V(G)$ are diffeomorphic to disks. In particular, the unbounded face is not in $V(G)$.
- (ii) If f_1, f_2 and f_3 are three faces of w which share one vertex, then at least one of the three does not belong to $V(G)$.
- (iii) If f_1 and f_2 belongs to $V(G)$ then every edge of $D(w)$ between f_1 and f_2 belongs to $E(G)$, ie G is an induced subgraph of $D(w)$.

If f is a vertex of G we define $\text{ed}(f)$, the external degree of f , by the formula

$$\text{ed}(f) = \text{deg}_{D(w)}(f) - 2 \text{deg}_G(f).$$

Remark Note that the external degree of a face f is always an even number because, w being bipartite, all cycles are of even length and hence $\text{deg}_{D(w)}$ is even. Condition (ii) ensures that $\text{ed}(f)$ is nonnegative.




Figure 9: Example of a red graph




Figure 10: Interpretation of the external degree in terms of gray half-edges. On the left is a portion of a web w with a red graph; on the right is the same portion of w with the vertices of G colored orange. The external degree of f is the number of half-edges adjacent to f which are not orange. In our case $\text{ed}(f) = 2$.

Let G be a red graph for w . If on the web we color the faces which belong to $V(G)$, then the external degree of a face f in $V(G)$ is the number of half-edges of w which are adjacent to a vertex of a cycle of w bounding f and lie in the uncolored region. These half-edges are called the *gray half-edges* of f in G , or of G when we consider the set of all gray half-edges of all vertices of G . See Figure 10.

An oriented red graph is a red graph together with an orientation. *A priori*, there is no restriction on the orientations, but as we shall see just a few of them will be relevant.

Definition 3.2 Let w be an ε -web, G be a red graph for w and o an orientation for G , we define the *level* $i_o(f)$ (or $i(f)$ when this is not ambiguous) of a vertex f of G by the formula

$$\begin{aligned} i_o(f) &\stackrel{\text{def}}{=} 2 - \frac{1}{2} \text{ed}(f) - \#\{\text{edges of } G \text{ pointing to } f\} \\ &= 2 - \frac{1}{2} \text{deg}_{D(w)}(f) + \#\{\text{edges of } G \text{ pointing away from } f\} \end{aligned}$$

and the *level* $I(G)$ of G is the sum of levels of all vertices of G .

The level is designed to locally encode the degree of a certain foam which we will show to be idempotent. The 2 corresponds to the opposite of the degree of a cup, while the other terms comes from zip moves (which have degrees equal to $+1$), this will become clear in [Lemma 3.18](#).

Remark The level is an integer because of the remark on page [1316](#). Note that the level of G does not depend on the orientation of G and

$$I(G) = 2\#V(G) - \#E(G) - \frac{1}{2} \sum_{f \in v(G)} \text{ed}(f).$$

Definition 3.3 A red graph is *admissible* if one can choose an orientation such that for each vertex f of G we have $i(f) \geq 0$. Such an orientation is called a *fitting orientation*. An admissible red graph G for w is *exact* if $I(G) = 0$.

Definition 3.4 Let w be an ε -web and G be a red graph for w . A *pairing* of G is a partition of the gray half-edges of G into subsets of 2 gray half-edges adjacent to the same face of w , one pointing towards it, and the other pointing away from it. A red graph together with a pairing is called a *paired red graph*.

Definition 3.5 A red graph G in an ε -web w is *fair* (resp. *nice*) if for every vertex f of G we have $\text{ed}(f) \leq 4$ (resp. $\text{ed}(f) \leq 2$).

Lemma 3.6 *If G is an admissible red graph in an ε -web w , then G is fair.*

Proof This follows directly from the definition of the level. □

Corollary 3.7 *Let w be a nonelliptic ε -web. If G is an admissible red graph for w then it has at least two edges.*

Proof If G were to contain just one vertex f , it would have external degree greater than or equal to 6, contradicting [Lemma 3.6](#). We can actually show that such a red graph contains at least 6 vertices (see [Corollary 4.12](#) and [Proposition 4.15](#)). □

Remark If a red graph G is nice, there is only one possible pairing. If it is fair, the number of pairings is 2^n , where n denotes the number of vertices with external degree equal to 4.

If in a picture one draws together a web w and a red graph G for w , they can encode a pairing of G in the picture by joining⁴ with dashed lines the paired half-edges. Note that if G is fair it is always possible to draw disjoint dashed lines (see [Figure 11](#) for an example).




Figure 11: A web w , a red graph G and the two possible pairings for G

Definition 3.8 Let ε be an admissible sequence of signs of length n , an ε -web w is said to be *virtually indecomposable* if $\langle \bar{w}w \rangle$ is a monic symmetric Laurent polynomial of degree n . An ε -web which is not virtually indecomposable is *virtually decomposable*. If w is a virtually decomposable ε -web, we define the *level* of w to be the integer $\frac{1}{2}(\deg(\bar{w}w) - n)$. Note that, despite of its fractional definition, the level is an integer.

The rest of the paper is devoted to showing the following theorem.

Theorem 3.9 Let ε be an admissible sequence of signs of length n and w an ε -web. Then the K^ε -module P_w is indecomposable if and only if w is virtually indecomposable.

Remark We do not suppose that w is nonelliptic, but as a matter of fact, if w is elliptic then $\langle \bar{w}w \rangle$ is not monic of degree n and the module P_w is decomposable.

One direction of the equivalence is given by Lemma 2.17. To prove the other direction of Theorem 3.9 we use red graphs and we will show the following two statements; the first is proved in this section, the latter in Section 4.

Theorem 3.10 To every exact paired red graph of w we can associate a nontrivial idempotent of $\text{HOM}_{K^\varepsilon}(P_w, P_w)$. Furthermore the direct factor associated with this idempotent is isomorphic to a web module.

Theorem 3.11 If w is a nonelliptic virtually decomposable ε -web of level k , then w contains a red graph of level k .

From the existence of a red graph of level k in the web w alone, we cannot deduce immediately that P_w is decomposable. However, an easy study of the combinatorics of red graphs (Section 4B) will be enough to make this conclusion. Indeed, it follows immediately from Propositions 4.6 and 4.8 that when there exists a red graph in w of nonnegative level, there exists an exact red graph in w . Hence, Theorem 3.9 follows from Theorems 3.10 and 3.11.

⁴We impose that w intersect the dashed lines only at their ends.




Figure 12: Example of a G -reduction of an ε -web w . The dashed lines represent the pairing.

Definition 3.12 Let w be an ε -web and G a paired red graph for w . We define the G -reduction of w to be the ε -web denoted by w_G and constructed as follows (see Figure 12 for an example):

- (1) For every face of w which belongs (as a vertex) to G , remove all edges adjacent to this face.
- (2) For every face of w which belongs (as a vertex) to G , connect the gray half-edges of G according to the pairing.

Note that even when w is nonelliptic, w_G needs not to be nonelliptic.

Definition 3.13 Let w be an ε -web and G a fair paired red graph for w . We define the projection associated with G to be the (w, w_G) -foam denoted by p_G and constructed as follows (from bottom to top):

- (1) For every edge of G , unzip (see Figure 13) the corresponding edge in w . Note that condition (ii) in the definition of red graph implies that all these unzip moves are far from each other, therefore we can perform the unzips simultaneously.⁵ Let us denote by w' the ε -web at the top of the foam after this step. Each vertex of G corresponds canonically to a face of w' . Such a face is a circle, a digon or a square (with extra information given by the pairing) because G is fair, so every vertex of G has external degree smaller than or equal to 4.
- (2)
 - On each square of w' which corresponds to a vertex of G , perform a square move, following the pairing information, (see Figure 14).
 - On each digon of w' which corresponds to a vertex of G , perform a digon move (see Figure 14).
 - On each circle of w' which corresponds to a vertex of G , glue a cap (see Figure 14).

⁵We mean that the isotopy type of the resulting foam does not depend on the order in which the edges are unzipped.




Figure 13: Unzip on the edge e




Figure 14: A square move, a digon move and a cap. The same foams turned upside-down are called a reversed square move, a reversed digon move and a cup.

We also define i_G , the injection associated with G , to be the (w_G, w) -foam which is the mirror image of p_G with respect to the horizontal plane $\mathbb{R}^2 \times \{\frac{1}{2}\}$, and \tilde{e}_G to be the (w, w) -foam equal to $i_G \circ p_G$. An example can be seen in Figure 15.

Remark In the previous definition, the square, digons and circles of w' are two-by-two nonadjacent, hence the order in which one performs square, digon or cap moves does not change the isotopy class of the foam we construct.

Remark It's worthwhile to note that a digon move can be seen as an unzip followed by a cap, and that a square move can be seen as two unzips followed by a cap. With this point of view, we see that in i_G (and in p_G), to every edge of G and every pair of gray half-edges corresponds a zip (or an unzip) and every vertex of G corresponds a cup (or a cap).

Theorem 3.10 is an easy consequence of the following proposition.

Proposition 3.14 *If w is a nonelliptic web and G is an exact paired red graph for w , then the (w_G, w_G) -foam $p_G \circ i_G$ is equivalent under the foam relations of Definition 2.14 to a nonzero multiple of the identity (w_G, w_G) -foam $w_G \times [0, 1]$.*





Figure 15: At the top is a web together with a fair (actually nice) paired red graph G . At the bottom we see a movie representing e_G .

To prove this proposition we need to develop a framework in which to perform calculations with the explicit foams we gave in Definition 3.13.

3B Foam diagrams

Definition 3.15 Let w be an ε -web. A foam diagram κ for w consists of


- the ε -web w ,
- a fair paired red graph G ,
- a function δ (called a dot function for w) from $E(w)$, the set of edges of w , to \mathbb{N} . This function will be represented by the appropriate number of dots on each edge of w . The null function is denoted by 0 .

To a foam diagram κ we associate $f(\kappa)$, the (w_G, w_G) -foam given by $p_G \circ s_w(\delta) \circ i_G$, where $s_w(\delta)$ is the identity foam $\text{id}_w = w \times [0, 1]$ of w with exactly $\delta(e)$ dots on every facet $e \times [0, 1]$ (with $e \in E(w)$). The (w_G, w_G) -foam $f(\kappa)$ is equal to $p_G \circ i_G$, with dots encoded by δ . A foam diagram will be represented by the ε -web drawn together with the red graph, and with some dots added on the edges of the ε -web in order to encode δ .


We will often identify $\kappa = (w, G, \delta)$ with $f(\kappa)$ and it will be seen as an element of $\text{HOM}_{K^\varepsilon}(P_{w_G}, P_{w_G})$. We can rewrite some of the relations depicted in Figure 7 in terms of foam diagrams.

Proposition 3.16 *The relations depicted in Figure 7 give rises to the following relations on foams associated with foam diagrams:*


- The 3-dots relation




- The sphere relations




- The digon relations



- The square relations



- The E-relation



The dashed lines indicate the pairing. When the orientation of the ε -web is not depicted the relation holds for any orientation.




Figure 16: Extension of foam diagrams to partially oriented red graphs

Lemma 3.17 *Let w be an ε -web and $\kappa = (w, G, \delta)$ a foam diagram. Then $f(\kappa)$ is equivalent to a \mathbb{Z} -linear combination of $s_{w_G}(\delta_i) = f((w_G, \emptyset, \delta_i))$ for some dot functions δ_i for w_G .*

Proof Thanks to the E-relation of Proposition 3.16 one can express $f(\kappa)$ as a \mathbb{Z} -linear combination of $f((w_j, G_j, \delta_j))$, where the G_j are red graphs without any edge. Thanks to the sphere, the digon and square relations of Proposition 3.16, each $f((w_j, G_j, \delta_j))$ is equivalent either to 0 or to $\pm f(w_G, \emptyset, \delta'_j)$. □

Lemma 3.18 *Let w be an ε -web and $\kappa = (w, G, 0)$ a foam diagram, with G exact. Then $f(\kappa)$ is equivalent to a multiple of $w_G \times [0, 1]$.*

Proof From the previous lemma we know that $f(\kappa)$ is equivalent to a \mathbb{Z} -linear combination of $w_G \times [0, 1]$ with some dots on it. We will see that the foam $f(\kappa)$ has the same degree as the foam $w_G \times [0, 1]$. This will prove the lemma because adding a dot on a foam increases its degree by 2.

To compute the degree of $f(\kappa)$ we view it as a composition of elementary foams thanks to its definition:

$$\begin{aligned} \deg(f(\kappa)) &= \deg(w \times [0, 1]) + 2(-2\#V(G) + (\#E(G) + \#\{\text{gray half-edges of } G\}/2)) \\ &= l(\varepsilon) + 2 \cdot 0 \\ &= \deg(w_G \times [0, 1]). \end{aligned}$$

The first equality is due to the decomposition pointed out in the remark on page 1321 and because an unzip (or a zip) has degree +1 and a cap (or a cup) has degree -2. The factor 2 is due to the fact $f(\kappa)$ is the composition of i_G and p_G . The second equality follows from the exactness of G . □

To prove Proposition 3.14, we only need to show that in the situation of the last lemma, the multiple is not equal to zero. In order to evaluate this multiple, we extend foam diagrams to (partially) oriented fair paired red graphs by the local relation indicated in Figure 16.





Figure 17: A positive edge (left) and a negative edge (right).

By *partially oriented* we mean that some edges are oriented and some are not. If G is partially oriented, and κ is a foam diagram with red graph G , we say that κ' is the *classical foam diagram* associated with κ if it is obtained from κ by applying the relation of Figure 16 to every oriented edge. Note that κ and κ' represent the same foam.

Definition 3.19 If w is an ε -web, G a red graph for w and o a partial orientation of G , we define $\gamma(o)$ to be equal to $\#\{\text{negative edges of } G\}$. An oriented edge is called *positive* or *negative* according to the rule depicted in Figure 17.

Lemma 3.20 Let w be an ε -web and G a partially oriented red graph with a nonoriented edge e of G . Then we have the following equality of foams:



If G is an unoriented red graph for w , and δ a dot function for w , then

$$f(w, G, \delta) = \sum_o (-1)^{\gamma(o)} f(w, G_o, \delta),$$

where G_o stands for G endowed with the orientation o , and o runs through all the $2^{\#E(G)}$ complete orientations of G .

Proof The first equality is the translation of the E-relation (see Proposition 3.16) in terms of foam diagrams of partially oriented red graphs. The second formula is the expansion of the first one to all edges of G . □

Lemma 3.21 If w is an ε -web, G an exact paired red graph for w and o a nonfitting orientation for G , then the (w_g, w_G) -foam $f(w, G_o, 0)$ is equivalent to 0.





Figure 18: The 5 different local situations of a foam diagram κ' in a neighborhood of a vertex of G' . On the second line, the digon on the left is *positive* and the digon on the right is *negative*.

Proof The orientation o is a nonfitting orientation. Hence, there is at least one vertex v of G so $i_o(v) > 0$. There are two different situations, either $i_o(v) = 1$ or $i_o(v) = 2$. Using the definition of foam diagrams for oriented red graphs (Figure 16), we deduce that κ' the classical foam diagram associated with $f(w, G_o, 0)$ looks around v like one of the three following situations:



The sphere relations and the digon relations provided by Proposition 3.16 imply that the foam $f(w, G_o, 0)$ is equivalent to 0. □

Lemma 3.22 *If w is an ε -web, G an exact paired red graph for w and o a fitting orientation for G , then the (w_G, w_G) -foam $f(w, G_o, 0)$ is equivalent to $(-1)^{\mu(o)} w_G \times I$, where $\mu(o) = \#V(G) + \#\{\text{positive digons of } G_o\}$ (see definition in Figure 18).*

Proof Let $\kappa' = (w', G', \delta')$ be the classical foam diagram associated with $(w, G_o, 0)$. The red graph G' has no edge. Locally, the foam diagram κ' corresponds to one of the five situations depicted in Figure 18.

By using some of relations of Proposition 3.16 we can remove all the vertices of G' . We see that $f(w, G_o, 0)$ is equivalent to $(-1)^{\#V(G') - \#\{\text{positive digons}\}} w_G \times I$ because the positive digon is the only one with no minus sign in the relations of Proposition 3.16. This proves the result because $V(G) = V(G')$. □

Lemma 3.23 *If w is an ε -web, G an exact paired red graph for w , and o_1 and o_2 two fitting orientations for G , then $\mu(o_1) + \gamma(o_1) = \mu(o_2) + \gamma(o_2)$.*

Proof We consider the two classical foam diagrams $\kappa'_1 = (w', G', \delta_1)$ and $\kappa'_2 = (w', G', \delta_2)$ corresponding to $(w, G_{o_1}, 0)$ and $(w, G_{o_2}, 0)$, respectively.

The red graph G' has no edge, and the possible local situations are depicted in [Figure 18](#). Consider a vertex v of G' , then a side of the face of w corresponding to v is either clockwise or counterclockwise oriented (with respect to this face). From the definition of γ we obtain that for $i = 1, 2$, $\gamma(o_i)$ is equal to the number of dots in κ'_i on clockwise oriented edges in w' . The dot functions δ_1 and δ_2 differ only on the edges adjacent to the digons, so $\gamma(o_1) - \gamma(o_2)$ is equal to the number of negative digons in κ'_1 minus the number of negative digons in κ'_2 . So

$$\begin{aligned} \gamma(o_1) - \gamma(o_2) &= \mu(o_2) - \mu(o_1), \\ \gamma(o_1) + \mu(o_1) &= \gamma(o_2) + \mu(o_2). \end{aligned} \quad \square$$

Proof of Proposition 3.14 The foam $p_G \circ i_G$ is equal to $f(w, G, 0)$. From Lemmas [3.20–3.22](#) we have

$$\begin{aligned} f(w, G, 0) &= \sum_{o \text{ fitting orientation of } G} (-1)^{\gamma(o)} f(w, G_o, \delta) \\ &= \sum_{o \text{ fitting orientation of } G} (-1)^{\gamma(o) + \mu(o)} w_G \times [0, 1] \\ &= \pm \#\{\text{fitting orientations of } G\} (w_G \times [0, 1]). \end{aligned}$$

The red graph G was supposed to be exact. In particular, the set of fitting orientations is not empty so $p_G \circ i_G$ is a nontrivial multiple of $\text{id}_{w_G} = w_G \times [0, 1]$. \square

Proof of Theorem 3.10 From [Proposition 3.14](#), we know that there exists a nonzero integer λ_G such that $p_G \circ i_G = \lambda_G w_G \times [0, 1]$. Hence, $\frac{1}{\lambda_G} i_G \circ p_G$ is an idempotent. It is clear that it is nonzero, and quite reasonable that it is not equivalent to the identity foam; for a proper proof, see [Proposition 3.25](#) or the remark on page [1314](#). \square

Remark To construct p_G we need to invert λ_G . This is the reason we work over \mathbb{Q} . It would be interesting to know which values λ_G can take for a minimal red graph G , to know if we can work over other fields (or rings).

3C On the identity foam

Definition 3.24 Let w be an ε -web with no circle and f a (w, w) -foam. We say that f is *reduced* if every facet of f is diffeomorphic to a disk and if f contains no singular circle (ie only singular arcs). In particular this implies that every facet of f meets $w \times \{0\}$ or $w \times \{1\}$.

The aim of this section is to prove the following proposition.

Proposition 3.25 *Let w be a nonelliptic ε -web. If f is a reduced (w, w) -foam which is equivalent (under the foam relations of [Definition 2.14](#)) to a nonzero multiple of $w \times [0, 1]$, then the underlying prefoam is diffeomorphic to $w \times [0, 1]$ and contains no dot.*

For this purpose we begin with a few technical lemmas.

Lemma 3.26 *Let w be a closed web and e an edge of w . Then there exists a (\emptyset, w) -foam f which is not equivalent to 0 such that the facet of f adjacent to the edge e contains at least one dot.*

Proof We prove the lemma by induction on the number of edges of the web w . It is enough to consider the case in which w is connected because the functor \mathcal{F} is monoidal. If the web w is a circle this is clear, since a cap with one dot on it is not equivalent to 0. If w is the theta web, then this is clear as well, since the half theta foam with one dot on the facet meeting e is not equivalent to 0.

Else, there exists a square or digon in w disjoint from e . Let us denote by w' the web similar to w but with the digon replaced by a single edge or the square smoothed in one way or the other. By induction we can find an (\emptyset, w') -foam f' nonequivalent to 0 with one dot on the facet adjacent to e .

In a neighborhood of the strand or the smoothed square, we consider a digon move or a square move (move upside down the pictures of [Figure 14](#)). Seen as a (w', w) -foam it induces an injective map. Therefore, the composition of f' with this (w', w) -foam is not equivalent to 0 and has one dot on the facet adjacent to e . \square

Notation 3.27 Let w be an ε -web and e an edge of w . We denote by $f(w, e)$ the $(\emptyset, \bar{w}w)$ -foam which is diffeomorphic to $w \times [0, 1]$ with one dot on the facet $e \times [0, 1]$. We denote by $f(w, \emptyset)$ the $(\emptyset, \bar{w}w)$ -foam which is diffeomorphic to $w \times [0, 1]$ with no dot on it.

Corollary 3.28 *Let w be an ε -web, and e an edge of w . The foam $f(w, e)$ is nonequivalent to 0.*

Proof From [Lemma 3.26](#) we know that for any w , there exists a (w, w) -foam which is nonequivalent to 0 and is the product of $f(w, e)$ with another (w, w) -foam. This proves that the (w, w) -foam $f(w, e)$ is not equivalent to 0. \square




Figure 19: From left to right a \cap , a λ and an H

Definition 3.29 If w is an ε -web. We say that it contains a λ , a \cap or an H if in a neighborhood of the boundary, w looks like one of the pictures of Figure 19.

Lemma 3.30 Every nonelliptic ε -web contains at least a λ , a \cap or an H .

Proof The closed web $\bar{w}w$ contains a circle a digon or a square, and this happens only if w contains a \cap , λ or H . □

Remark In fact, one can “build” every nonelliptic web with these three elementary webs. This is done via the so-called growth algorithm [6].

Lemma 3.31 Let w be a nonelliptic ε -web. Then the elements of $(f(w, e))_{e \in E(w)}$ are pairwise nonequivalent (but they may be linearly dependent).

Sketch of the proof We proceed by induction on the number of edges of w . The base case is straightforward, since if w has only one edge there is nothing to prove. We can distinguish several case thanks to Lemma 3.30.

If w contains a \cap , we denote by e the edge of this \cap , and by w' the ε' -web similar to w but with the \cap removed. Suppose that $e_1 = e$. Then $e_2 \neq e$ and the $(\emptyset, \bar{w}w)$ -foams $f(w, e_1)$ and $f(w, e_2)$ are nonequivalent because if we glue a cap with one dot on it on the cup (we mean $e \times I$), on the one hand we obtain a $(\emptyset, \bar{w}'w')$ -foam equivalent to 0 and on the other hand a $(\emptyset, \bar{w}'w')$ -foam equivalent to $-f(w', \emptyset)$. Thanks to Lemma 3.26, we know that this last $(\emptyset, \bar{w}'w')$ -foam is not equivalent to 0. If e_1 and e_2 are different from e , it is clear as well, because $f(w, e_1)$ and $f(w, e_2)$ can be seen as composite of $f(w', e_1)$ and $f(w', e_2)$ with a birth (seen as a $(\bar{w}'w', \bar{w}w)$ -foam) which is known to correspond to injective map.

The same kind of argument shows the two other cases, but the digon relations and the square relations are used instead of the sphere relations. □

Lemma 3.32 Let w be an ε -web with no closed connected component (ie w is ∂ -connected; see Section 4C) and f a reduced (w, w) -foam f . Suppose that every facet meets $w \times \{0\}$ on at most one edge, and meets $w \times \{1\}$ on at most one edge. Then it is isotopic to $w \times [0, 1]$.

Proof The proof is inductive on the number of vertices of w . If w is a collection of arcs, the foam f has no singular arc. As f was supposed to be reduced, it has no singular circle. Therefore it is a collection of disks which corresponds to the arcs of w , which proves the result in this case.

We suppose now that w has at least one vertex. Let us pick a vertex v which is a neighbor (via an edge e) of the boundary ε of w . We claim that the singular arc α starting at $v \times \{0\}$ must end at $v \times \{1\}$.

Indeed, the arc α cannot end at $w \times \{0\}$, for otherwise, the facet f adjacent to e would be adjacent to another edge of w . Therefore the arc α ends at $w \times \{1\}$. For exactly the same reasons, it has to end at $v \times \{1\}$, so the facet which is adjacent to $e \times \{0\}$ is isotopic to $e \times I$. Now we can remove a neighborhood of this facet and we are back in the same situation with a ε' -web with fewer vertices, and this concludes the proof. \square

Proof of Proposition 3.25 We consider a nonelliptic ε -web w . Let f be a reduced (w, w) -foam such that f is equivalent to $w \times I$ up to a nontrivial scalar. We claim that the foam f satisfies the hypotheses of Lemma 3.32. First, the web w , being nonelliptic, has no closed connected component. Suppose that a facet of f meets $w \times \{0\}$ at two different edges e_1 and e_2 . The foams obtained by composing f with $f(w, e_1)$ and $f(w, e_2)$ would be equal. This contradicts Lemma 3.31. The same argument shows that every facet meets $w \times \{1\}$ on at most one edge. Hence, the foam f does satisfy the hypotheses of Lemma 3.32, and f is isotopic to $w \times [0, 1]$. \square

We conjecture that Proposition 3.25 still holds without the nonellipticity hypothesis. However the proof has to be changed since Lemma 3.31 and Lemma 3.32 cannot be extended to elliptic webs.

Corollary 3.33 *If w is a nonelliptic ε -web and w' is an ε -web with strictly fewer vertices than w , then if f is a (w, w') -foam and g is a (w', w) -foam, the (w, w) -foam fg is not equivalent to a scalar times the identity.*

4 Characterization of indecomposable web modules

4A General view

The purpose in this section is to prove Theorem 3.11. The proof is an induction on the number of edges of the web w . But for the induction to work, we need to handle elliptic webs and to introduce *stacks of red graphs* and ∂ -connectedness (Section 4C).

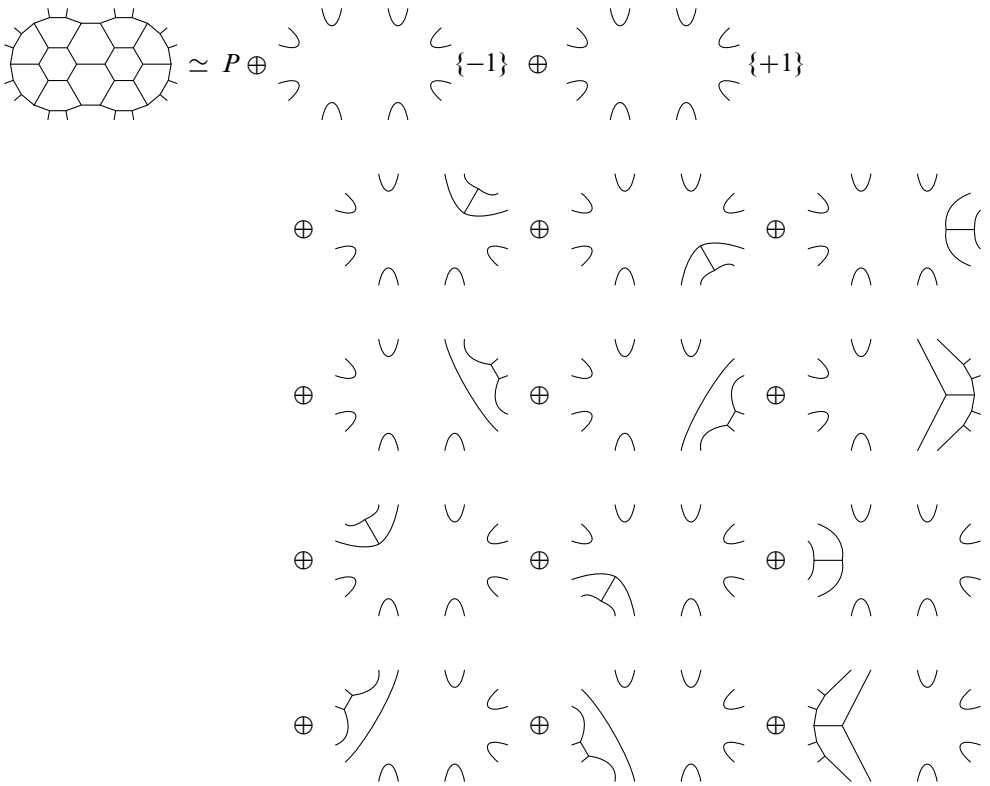


Figure 20: Example of a decomposition of a web module into indecomposable modules. All direct factors which are web modules are obtained through idempotents associated with red graphs. The module P is not a web module but is a projective indecomposable module.

Definition 4.1 Let w be an ε -web. A stack of red graphs $S = (G_1, G_2, \dots, G_l)$ for w is a finite sequence of paired red graphs such that G_1 is a red graph of $w_1 \stackrel{\text{def}}{=} w$, G_2 is a red graph of $w_2 \stackrel{\text{def}}{=} w_{G_1}$, G_3 is a red graph of $w_3 \stackrel{\text{def}}{=} (w_{G_1})_{G_2} = (w_2)_{G_2}$ etc. We denote $(\dots((w_{G_1})_{G_2})\dots)_{G_l}$ by w_S and we denote l by $l(S)$ and say that it is the length of S . We define the level of a stack to be the sum of the levels of the red graphs of the stack.

Definition 4.2 A stack of red graphs is nice if all its red graphs are nice. Note that in this case the pairing information on red graphs is empty.

Definition 4.3 An ε -web w is ∂ -connected if every connected component of w meets the boundary.

A direct consequence is that a ∂ -connected ε -web contains no circle.




Figure 21: A stack of red graphs of length 2

Lemma 4.4 *A nonelliptic ε -web is ∂ -connected.*

Proof An ε -web which is not ∂ -connected has a closed connected component. This connected component contains at least a circle, digon or square and hence is elliptic. \square

We can now rewrite [Theorem 3.11](#) as it will be proven.

- Proposition 4.5** (1) *If w is a ∂ -connected ε -web which is virtually decomposable of level $k \geq 1$ then there exists a nice stack of red graphs S for w of level greater than or equal to k such that w_S is ∂ -connected.*
- (2) *If w is a ∂ -connected ε -web which is virtually decomposable of level $k \geq 1$, contains no digon and contains exactly one square which is supposed to be adjacent to the unbounded face then there exists a nice red graph G in w of level greater than or equal to k such that w_G is ∂ -connected.*
- (3) *If w is a nonelliptic ε -web which is virtually decomposable of level $k \geq 0$ then there exists a nice red graph G in w of level greater of equal to k such that w_G is ∂ -connected.*

We will prove [Proposition 4.5](#) in [Section 4D](#), thanks to a technical lemma ([Lemma 4.22](#)) which will be proven in [Section 4F](#) after an alternative perspective on red graphs ([Section 4E](#)). Before this we discuss combinatorics of red graphs and ∂ -connectedness.

Remark It is easy to see that a nonelliptic superficial ε -web contains no red graphs of nonnegative level, hence this result is strictly stronger than an earlier result of the author [[14](#)].

4B Combinatorics on red graphs

In our construction, it is very difficult to ensure exactness or even admissibility of the red graph. The aim of this section is to show that all we need to focus on is the index of the red graphs we build. Proposition 4.6 tells that whenever we find a red graph in a web with nonnegative index, we can find an admissible red graph with nonnegative index, and Proposition 4.8 indicates that if we can find an admissible red graph in a web, then there exists an exact red graph for this web.

Proposition 4.6 *Let w be an ε -web. Suppose there exists a red graph G for w such that $I(G) \geq 0$; then there exists an admissible red graph \tilde{G} for w such that $I(\tilde{G}) \geq I(G)$.*

Proof If G is already admissible, there is nothing to show; hence we suppose that G is not admissible. Among all the orientations for G , we choose one such that $\sum_{f \in V(G)} |i(f)|$ is minimal and we denote it by o . From now on G is endowed with this orientation. As G is not admissible there exist some vertices with negative level and some with positive level.

We first show that there is no oriented path from a vertex f_p with $i_o(f_p) > 0$ to a vertex f_n with $i_o(f_n) < 0$. Suppose there exists such a path and denote it by γ . Let us inspect the orientation o' , which is the same as o except along the path γ where it is reversed. For all vertices f of G but f_p and f_n , we have $i_o(f) = i_{o'}(f)$ and

$$i_{o'}(f_p) = i_o(f_p) - 1, \quad i_{o'}(f_n) = i_o(f_n) + 1.$$

But then $\sum_{f \in V(G)} |i_{o'}(f)|$ would be strictly smaller than $\sum_{f \in V(G)} |i_o(f)|$ and this contradicts that o is minimal.

We consider the induced oriented subgraph (\tilde{G}, \tilde{o}) of (G, o) whose vertices are exactly those of G which can be reached from a vertex with positive level by an oriented path. This set is not empty since it contains the vertices with positive degree. It contains no vertex with negative degree. For all vertices of \tilde{G} , we have

$$\begin{aligned} i_{\tilde{o}}(f) &= 2 - \frac{1}{2} \deg_{D(w)}(f) + \#\{\text{edges of } \tilde{G} \text{ pointing away from } f\} \\ &= 2 - \frac{1}{2} \deg_{D(w)}(f) + \#\{\text{edges of } G \text{ pointing away from } f\} \\ &= i_o(f). \end{aligned}$$

The second equality holds because if f is in $V(\tilde{G})$, every edge in $E(G) \setminus E(\tilde{G})$ points to f by definition of \tilde{G} . This shows that \tilde{G} is admissible and $I(\tilde{G}) > I(G)$. \square

Lemma 4.7 *Let w be a nonelliptic web, suppose that it contains a red graph of level k . Then it contains an admissible nice red graph of level at least k .*

Proof We consider a red graph G of w of level k . Thanks to [Proposition 4.6](#) we can suppose that it is admissible. We can take a minimal red graph G (for the inclusion of the set of vertices) for the property of being of level at least k and admissible. The graph G is endowed with a fitting orientation. Now suppose that it is not nice, it means that there exists a vertex v of G which has external degree equal to 4. But G being admissible all the edges of G adjacent to v point out of v , so we can remove v ; ie we can consider the induced subgraph G' with all the vertices of G but v and with the induced orientation. Then the red graph G' is admissible with the same level. Hence G is not minimal, which is a contradiction. \square

For a nonelliptic ε -web, the existence of an exact red graph may appear as an exceptional situation between the case where there is no admissible red graph and the case where all admissible red graphs are nonexact. The aim of the rest of this section is to show [Proposition 4.8](#), which indicates that this is not the case. Along the way we state some small results that are not directly useful for the proof but may enlighten us as to what red graphs look like.

Proposition 4.8 *Let w be a nonelliptic ε -web. If there exists an admissible red graph for w then there exists an exact red graph for w .*

Definition 4.9 Let w be an ε -web, and G and G' two admissible red graphs for w . We say that G' is a *red subgraph* of G if $V(G') \subset V(G)$. We denote by $\mathcal{G}(G)$ the set of all admissible red subgraphs. It is endowed with the partial order given by the inclusion of sets of vertices. We say that G is *minimal* if $\mathcal{G}(G) = \{G\}$.

Note that a red subgraph is an induced subgraph and that a minimal red graph is connected.

Lemma 4.10 *Let w be an ε -web and G a minimal admissible red graph endowed with a fitting orientation. Then there is no nontrivial partition of $V(G)$ into two sets V_1 and V_2 such that for each vertex v_1 in V_1 and each vertex v_2 in V_2 every edge between v_1 and v_2 is oriented from v_1 to v_2 .*

Proof If there were a such a partition, we could consider the red subgraph G' with $V(G') = V_2$. For every vertex in V_2 , the level is the same in G and G' and, hence, G' would be admissible and G would not be minimal. \square

Corollary 4.11 *Let w be an ε -web and G a minimal admissible red graph for w . Then the graph G has no leaf.⁶ Therefore if it has 2 or more vertices, it is not a tree.*

Proof Let us endow G with a fitting orientation. If v were a leaf of G , the vertex v would be either a sink or a source, hence $V(G) \setminus \{v\}$ and $\{v\}$ would partition $V(G)$ in a way forbidden by [Lemma 4.10](#). □

Corollary 4.12 *If G is an admissible red graph for a nonelliptic ε -web w , then G is not a tree.*

Proof Consider a minimal red subgraph of G . Thanks to [Corollaries 3.7](#) and [4.11](#), it is not a tree. Hence G is not a tree. □

Lemma 4.13 *Let w be an ε -web and G a minimal admissible red graph for w . If G has more than 2 vertices, then it is nice.*

Proof Suppose that we have a vertex v of G with external degree equal to 4. Consider a fitting orientation for G . All edges of G adjacent to v would point out, otherwise the degree of v would be negative. So v would be a sink and, thanks to [Lemma 4.10](#), this is not possible. □

Lemma 4.14 *Let w be a nonelliptic ε -web and G a minimal admissible red graph. If the red graph G is endowed with a fitting orientation, then it is strongly connected.*

The terms *weakly connected* and *strongly connected* are classical in graph theory. The first means that the underlying unoriented graph is connected in the usual sense. The second that for any pair of vertices (v_1, v_2) , there exists an oriented path from v_1 to v_2 and an oriented path from v_2 to v_1 .

Proof Let v be a vertex of G , consider the subset V_v of $V(G)$ which contains the vertices of G reachable from v by an oriented path. The sets V_v and $V(G) \setminus V_v$ form a partition of $V(G)$ which must be trivial because of [Lemma 4.10](#), but v is in V_v so $V_v = V(G)$. This is true for any vertex v , so G is strongly connected. □

Proposition 4.15 *If G is a red graph for a nonelliptic ε -web w , then any (not oriented) simple cycle has at least 6 vertices.*

Proof Take a nontrivial simple cycle C in G . We consider the collection of faces of w nested inside C (this is nonempty thanks to condition (iii) of the definition of red graphs). This defines a plane⁷ graph H . We define H' to be the graph H with the bivalent vertices smoothed (we mean here that if locally H looks like $\bullet\text{---}\bullet$, then H' looks like $\text{---}\text{---}$). An example of this construction is depicted in [Figure 22](#).

⁶We mean vertex of degree 1.

⁷We mean a graph embedded in the plane without any crossing.




Figure 22: On the left the ε -web w and the red graph G , in the middle the graph H , and on the right the graph H'

The ε -web w being nonelliptic, each face of H has at least 6 sides. We compute the Euler characteristic of H' :

$$\chi(H') = \#F(H') - \#E(H') + \#V(H') = 2.$$

As in Proposition 2.2, this gives us $\sum_{i \in \mathbb{N}} \#F_i(H')(1 - i/6) = 2$, where $F_i(H')$ is the set of faces of H' with i sides. Restricting the sum to $i \leq 5$ and considering F'_i the set of bounded faces, we have

$$\sum_{i=0}^5 \#F'_i(H')(6 - i) \geq 6.$$

But the bounded faces of H' with fewer than 6 sides come from bounded faces of H which have at least 6 sides. The number n of bivalent vertices in H is therefore greater than or equal to $\sum_{i=0}^5 F'_i(H')(6 - i)$; ie greater than or equal to 6. But n is also the length of the cycle C . □

Note that a cycle in a red graph can have an odd length (as in the example of Figure 22).

Lemma 4.16 *Let G be a minimal admissible red graph for a nonelliptic ε -web w . Then G has at least one vertex with degree 2.*

Proof Suppose that all vertices of G have degree greater than or equal to 3, then the graph G would contain a face with less than 5 sides (this is the same argument as in Proposition 2.2 which states that a closed web contains a circle, a digon or a square). This contradicts Proposition 4.15. □

Proposition 4.8 is a direct consequence of the following lemma.




Figure 23: On the left w , on the right w'

Lemma 4.17 *Let w be a nonelliptic ε -web and G a minimal admissible red graph for w . Then G is exact.*

Proof We endow G with a fitting orientation o . Suppose G is not exact; then we can find a vertex f with $i_o(f) > 0$.

First consider the case where $\deg(f) = 2$. The ε -web w being nonelliptic, $\text{ed}(f) \geq 2$. This shows that the two edges adjacent to f point away from f . Hence, f is a sink and this contradicts Lemma 4.10.

Now, let us consider the general case. Let f' be a vertex with degree 2. Lemma 4.14 implies that there exists an oriented path γ from f to f' . Let us reverse the orientations of the edges of γ . We denote by o' this new orientation. Then we have $i_{o'}(f) = i_o(f) - 1 \geq 0$ and $i_{o'}(f') = i_o(f') + 1 \geq 1$. The levels of all other edges are not changed, hence o' is a fitting orientation, and we are back in the first situation (where f' plays the role of f). \square

4C The ∂ -connectedness

Lemma 4.18 *Let w be a ∂ -connected ε -web with a digon. The ε -web w' obtained from w by replacing a digon by a single edge (see Figure 23) is still ∂ -connected. In other words ∂ -connectedness is preserved by digon reduction.*

Proof This is clear because every path in w can be projected onto a path in w' . \square

Note that ∂ -connectedness is not preserved by square reduction, see for example Figure 24. However we have the following lemma.

Lemma 4.19 *If w is a ∂ -connected ε -web which contains a square S , then one of the two ε -webs obtained from w by a reduction of S (see Figure 25) is ∂ -connected.*

Proof Consider the oriented graph \tilde{w} obtained from w by removing the square S and the 4 half-edges adjacent to it (see Figure 26). We obtain a graph with 4 fewer trivalent vertices than w and 4 more univalent vertices than w . We call E_S the cyclically




Figure 24: The ∂ -connectedness is not preserved by square reduction.




Figure 25: The ε -web w with the square S (left), and the two reductions of the square S .




Figure 26: On the left w , on the right \tilde{w}

ordered set of the 4 new univalent vertices of \tilde{w} . The orientations of the vertices in E_S are $(+, -, +, -)$. Note that in \tilde{w} , the flow modulo 3 (see the remark on page 1305) is preserved everywhere, so the sum of the orientations of the univalent vertices of any connected component must be equal to 0 modulo 3. Suppose now that there is a connected component t of \tilde{w} which has all its univalent vertices in E_S . The flow condition implies that either all vertices of E_S are vertices of t or exactly two consecutive vertices of E_S are vertices of t , or that t has no univalent vertex. The first situation cannot happen because by adding the square S to t we would construct a closed connected component of w , which is supposed to be ∂ -connected. The last situation cannot happen either for the same reason. Hence the only situation that can occur is the second one. If there were two different connected components t_1 and t_2 of \tilde{w} such that t_1 and t_2 have all their vertices of degree 1 in E_S , then adding the square S to $t_1 \cup t_2$ would lead to a closed connected component of w . Hence, there is at most one connected component of \tilde{w} with all vertices of degree 1 in E_S . Call these vertices e_+ and e_- and call e'_+ and e'_- the two other vertices of E_S (the indices give the orientation). If we choose w' to be the ε -web corresponding to the smoothing which connects e_+ with e'_- and e_- with e'_+ , then w' is ∂ -connected. \square




Figure 27: On the left w_{S_0} , on the right w . If t_{S_0} is a circle, then w contains a digon.




Figure 28: On the left w_{S_0} , on the right w . If t_{S_0} contains a digon then w contains a square adjacent to S_0 .

Definition 4.20 Let w be a ∂ -connected ε -web and S a square in w . The square S is a ∂ -square if the two ε -webs $w_{=}$ and $w_{||}$ obtained from w by the two reductions of S are ∂ -connected.

Lemma 4.21 If w is a ∂ -connected web, then either it is nonelliptic, or it contains a digon or a ∂ -square.

Proof Suppose that w is not nonelliptic. As w is ∂ -connected, it contains no circle. It must contain at least a digon or a square. If it contains a digon we are done, so suppose w contains no digon. We should show that at least one square is a ∂ -square. Suppose that there is no ∂ -square. This means that for every square S , there is a reduction such that the resulting ε -web $w_{s(S)}$ obtained from w by a reduction of S has a closed connected component t_S . Let us consider a square S_0 such that t_{S_0} is as small as possible (in terms of the number of vertices, for example). The web t_{S_0} is closed and connected, so either it is a circle, or it contains a digon or at least two squares. If t_{S_0} is a circle then w contains a digon adjacent to the square S_0 , and we excluded this case (see [Figure 27](#)).

If it contains a digon, the digon must intersect any regular neighborhood of S_0 or else the digon would already be in w . Hence, it appears that the digon comes from a square S_1 in w (S_1 is adjacent to S_0), and t_{S_1} has two fewer vertices than t_{S_0} , which is excluded by the minimality of T_{S_0} ; see [Figure 28](#).




Figure 29: On the left w' , on the right w with the red graph G

Hence, the closed web t_{S_0} contains at least two squares. Let us pick a square S' which is disjoint from S_0 and hence comes from a square in w . At least one of the two smoothings of the square S' must disconnect t_{S_0} , else the square S' would be a ∂ -square in w . But as it disconnects t_{S_0} , $t_{S'}$ is a strict subgraph of t_{S_0} , and this contradicts the minimality of S_0 . This concludes the proof that w must contain a ∂ -square. \square

4D Proof of Proposition 4.5

In this section, we prove Proposition 4.5 admitting the following technical lemma which will be proven in Section 4F.

Lemma 4.22 *Let w be a ∂ -connected ε -web which contains no digon and at most one square which must be adjacent to the unbounded face. Let G be a nice red graph of w and G' a nice red graph of w_G such that w_G and $w_{G'}$ are ∂ -connected. Then there exists a red graph G'' of w such that $(w_G)_{G'} = w_{G''}$ and the level of G'' is greater than or equal to the sum of the levels of G and G' .*

This lemma says that under certain conditions one can “flatten” two red graphs.

Proof of Proposition 4.5 As we announced this will be done by induction on the number of edges of w . Let us suppose that (1), (2) and (3) hold for all ε -webs with strictly less than n edges, and let us consider an ε -web with n edges. Note that whenever w is nonelliptic, statement (3) is stronger than statement (1), so it suffices to prove (2) and (3) in this case.

We first prove (1). If w contains a digon, then we apply the induction hypothesis to w' , the ε -web similar to w but with the digon reduced (ie replaced by a single strand). The red graph G which consists of only one vertex (the digon) and no edge is nice and has level equal to 1 (see Figure 29).

If w' is not virtually decomposable or virtually decomposable of level 0, then w is virtually decomposable of level 1. In this case, the stack with only one red graph equal

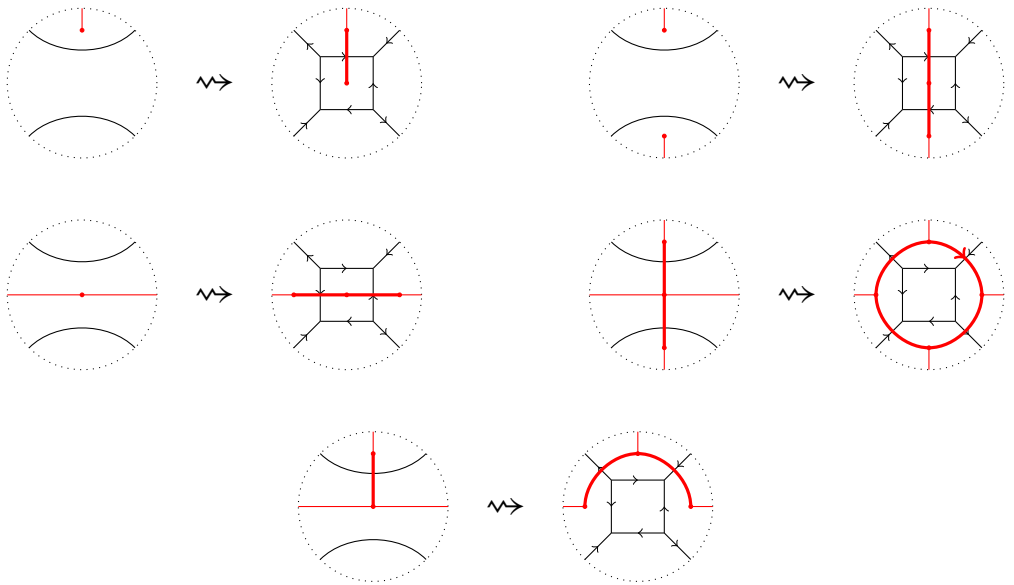


Figure 30: Transformations of G' to obtain G

to G is suitable and we are done. Else we know that w' is of a certain level $k - 1$ and that there exists a nice stack of red graphs S' of level at least $k - 1$ in w' and we consider the stack S equal to the concatenation of G with S' . This is a nice stack of red graphs of level at least k and we are done.

Suppose now that the ε -web w contains no digon, but a square, then it contains a ∂ -square (see Lemma 4.21). Suppose that the level of w is $k \geq 1$ (else there is nothing to show); then at least one of the two reductions is virtually decomposable of level k (this is a Cauchy–Schwartz inequality; see [15, Section 1.1] for details). Then we consider w' the ε -web obtained by a reduction of the square so that it is of level k . From the induction hypothesis we know that there exists a stack of red graphs S' in w' of level k . If all the red graphs of S' are disjoint from a neighborhood of the square, then we can transform the stack S' into a stack of w with the same level. Else, we consider the first red graph G' of S' which is not disjoint from the square location and according to the situation we define G by the moves given in Figure 30.

Replacing G' by G we can transform the stack S' into a stack for the ε -web w . The level of G is equal to the level of G' in all case but the second on the left. In this case, we have $I(G) = I(G') + 2$. Hence, in all cases, we have $I(G) \geq I(G')$, and therefore the level of S is greater than or equal to k .



Figure 31: The H of w (left) and its two reductions w_{\parallel} (middle) and w_{-} (right)

We now prove (2). From what we just did, we know that w contains a nice stack of red graphs of level k . Among all the nice stacks of red graphs of w with level greater than or equal to k , we choose one with minimal length and call it S . If its length were greater than or equal to 2, then Lemma 4.22 would tell us that we could take the first two red graphs and replace them by just one red graph with level greater than or equal to the sum of their two levels, so S would not be minimal. This proves that S has length 1, therefore, w must contain a nice red graph of level at least k .

We now prove (3). The boundary of w contains at least a \cap , a λ or an H (see Figure 19). In the first two cases, we can consider w' the ε -web with the \cap removed or the λ replaced by a single strand. Then w' is nonelliptic and virtually decomposable of level k and there exists a nice red graph in w' of level at least k . This red graph can be seen as a red graph of w , and we are done. If the boundary of w contains no λ and no \cap , then it must contain an H . There are two ways to reduce the H (see Figure 31). At least one of the two following situations happens: w_{\parallel} is virtually decomposable of level k or w_{-} is virtually decomposable of level $k + 1$.

In the first situation, we can use the same reasoning as before: w_{\parallel} being nonelliptic, the induction hypothesis gives that we can find a nice red graph G_{\parallel} of level at least k in w_{\parallel} . The red graph G_{\parallel} can be seen as a red graph of w , we denote it by G . The red graph G is nice and has level at least k and we are done. In the second situation, we consider w_{-} , to which we can apply the induction hypothesis (we are either in case (2) or in case (3)), so we can find a nice red graph G_{-} of level at least $k + 1$. The red graph G_{-} can be seen as a red graph of w , we denote it by G . The only possible difference between G and G_{-} is as follows: one vertex of G may have two more gray half-edges than the corresponding vertex of G_{-} . Hence the red graph G has level at least k but may not be nice. We can conclude the proposition thanks to Lemma 4.7. \square

4E A new approach to red graphs

In this section, we give an alternative approach to red graphs: instead of starting with a web and simplifying it with a red graph, we construct a red graph from a web and a simplification of this web. For this we need a property of webs that we have not used so far.




Figure 32: On the left a positive crossing, on the right a negative one. The path is dashed and the web is solid.

Proposition 4.23 *Let w be a closed web. Then it admits a (canonical) face 3-coloring with the unbounded face colored by an arbitrary but fixed color $c \in \mathbb{Z}/3\mathbb{Z}$. We call this coloring the face coloring of base c of w . When c is not mentioned it is meant to be 0.*

Proof We will color the connected components of $\mathbb{R}^2 \setminus w$ with elements of $\mathbb{Z}/3\mathbb{Z}$. We can consider the only unbounded component U of $\mathbb{R}^2 \setminus w$. We color it by c ; then for each other connected component f , we consider p an oriented path from a point in the interior of U to a point in the interior of f , which crosses w transversely. We then define the color of f to be the sum (modulo 3) of the signs of the points of intersection of the path p with w (see Figure 32 for sign conventions). This does not depend on the path because in w the flow is always preserved modulo 3 (see the remark on page 1305). And, by definition, two adjacent faces are separated by an edge, so they do not have the same color. □

Corollary 4.24 *Let w be an ε -web. Then the connected component of $\mathbb{R} \times \mathbb{R}_+ \setminus w$ admits a (canonical) 3-coloring with the unbounded connected component colored by c . We call this coloring the face coloring of base c of w .*

Proof We complete w with \bar{w} and use the previous proposition to obtain a coloring of the faces of \bar{w} . By restriction, this gives us a canonical coloring for $\mathbb{R} \times \mathbb{R}_+ \setminus w$. □

Note that in this corollary it is important to consider the connected component of $\mathbb{R} \times \mathbb{R}_+ \setminus w$ instead of the connected component of $\mathbb{R}^2 \setminus w$. Let us formalize this in a definition.

Definition 4.25 If w is an ε -web, the *regions* of w are the connected components of $\mathbb{R} \times \mathbb{R}_+ \setminus w$. The *faces* of w are the regions which do not intersect $\mathbb{R} \times \{0\}$.

Definition 4.26 Let w be an ε -web. An ε -web w' is a *simplification* of w if

- the set of vertices of w' is included in the set of vertices of w ,

- every edge e of w' is divided into an odd number of intervals $([a_i, a_{i+1}])_{i \in [0, 2k]}$ such that for every i in $[0, k]$, $[a_{2i}, a_{2i+1}]$ is an edge of w (with matching orientations) and for every i in $[0, k - 1]$, $]a_{2i+1}, a_{2i+2}[$ lies in the interior of the connected component of $\mathbb{R} \times \mathbb{R}_+ \setminus w$ opposite to $[a_{2i}, a_{2i+1}]$ with respect to a_{2i+1} (see Figure 33).




Figure 33: Local picture around a_j . The edge of w' is orange and large, while the ε -web w is black and thin.




Figure 34: The ε -web w (in black) and w_0 (in orange) of Proposition 2.19 seen in terms of simplification.

Note that in this definition the embedding of w' with respect to w is very important.

Lemma 4.27 *Let w be an ε -web and w' a ∂ -connected simplification of w . Suppose that an edge e of w is a (part of an) edge of w' as well. Then, for any color c , in the face colorings of base c of w and w' , the regions adjacent to e in w and in w' are colored in the same way.*

Proof If the edge e has one end on the boundary, this is obvious. Else, w' being ∂ -connected, we can build a path from the boundary to a point in a regular neighborhood of e following some edges of w' . This path intersects the web w more times than w' , but the signs of these extra intersections annihilate two by two, so, in the face colorings

of base c of w and w' , the regions adjacent to e in w and in w' are colored in the same way. □

Definition 4.28 Let w be an ε -web and w' a simplification of w . We consider the face colorings of w and w' . A face f of w lies in one or several regions of w' . This face f is *essential with respect to w'* if all regions of w' which it intersects have a different color than f .

Remark We could have given a more general definition for regions, but it is easy to see that a region of w which is not a face w is never essential.

Lemma 4.29 Let w be a ∂ -connected ε -web and w' a ∂ -connected simplification of w . If a face f of w is not essential with respect to w' then it intersects only one region of w' .

Proof Consider a face f of w which intersects more than one region of w' . We will prove that it is essential with respect to w' . Consider an edge e' of w' which intersects f (there is at least one by hypothesis). The intersection of e' with the boundary of f consists of some vertices of w (at least two). Let v be one of these vertices. A regular neighborhood of v is depicted in [Figure 35](#).




Figure 35: Above v the colors of w and w' coincide thanks to [Lemma 4.27](#).

We want to prove that none of the faces of w' which are adjacent to e' has the same color as the face f . This follows from the [Lemma 4.27](#), and from the fact that the part of e' above v is an edge of w (see [Figure 35](#)). □

Proposition 4.30 Let w be a ∂ -connected ε -web (this implies that every face of w is diffeomorphic to a disk) and w' a ∂ -connected simplification of w . Then there exists a (canonical) paired red graph G such that w' is equal to w_G . We denote it by $G_{w \rightarrow w'}$.

Proof We consider the canonical colorings of the faces of w and w' . The red graph G is the induced subgraph of $D(w)$ (the dual graph of w) whose vertices are essential faces of w with respect to w' . The pairing is given by the edges of w' . We need to




Figure 36: The three configurations for the vertex v of w : a vertex of w' (left), inside an edge of w' (middle), inside a region of w' (right).

prove first that this is indeed a red graph, and in a second step that $w_G = w'$. We consider a vertex v of w and the 3 regions adjacent to this vertex. We want to prove that at least one of the 3 regions is not essential with respect to w' .

There are three different situations (depicted in Figure 36):

- (a) The vertex v is a vertex of w' .
- (b) The vertex v is in the interior of an edge of w' .
- (c) The vertex v is in the interior of a face of w' .

In case (a), Lemma 4.27 proves that none of the three regions is essential.

Consider the case (b). One of the three regions intersects two different regions of w' , hence it is essential thanks to Lemma 4.29; the two others are not.

In case (c), the 3 regions have different colors so that one of them has the same color as the region of w' in the interior of which v lies. This region is therefore not essential. This shows that G is a red graph (we said nothing about the admissibility).

Let us now show that $w' = w_G$. We consider a collection $(N_f)_{f \in \mathcal{V}(G)}$ of regular neighborhoods of the essential faces of w with respect to w' . Let us first show that for every essential face f of w , the restrictions to N_f of w_G and w' match. As f is essential, it is a vertex of G . Then the restriction of w_G to N_f is a collection of strands joining different points of the boundary according to the pairing, just as w' .

In $\mathbb{R} \times \mathbb{R}_+ \setminus (\bigcup_{f \in \mathcal{V}(G)} N_f)$ the ε -webs w' and w_G are both equal to w . This completes the proof. □

Note that $G_{w \rightarrow w'}$ depends on how w' is embedded to see it as a simplification of w .

Definition 4.31 Let w an ε -web and w' a simplification of w . Then the simplification is *nice*, if for every region r of w , $r \cap w'$ is either the empty set or connected.

This leads to a natural lemma.

Lemma 4.32 *Let w be a ∂ -connected ε -web and w' a ∂ -connected simplification of w . The simplification is nice if and only if the red graph $G_{w \rightarrow w'}$ is nice*

Proof Thanks to Lemma 4.29, only essential faces of w with respect to w' can have nontrivial intersection with w' , and for an essential face f , twice the number of connected components of $f \cap w'$ is equal to the external degree of the vertex of $G_{w \rightarrow w'}$ corresponding to f . □

Lemma 4.33 *If w is a ∂ -connected ε -web, and w' a ∂ -connected simplification of w . Then the level of $G_{w \rightarrow w'}$ is given by*

$$I(G_{w \rightarrow w'}) = 2 \#\{\text{essential faces of } w \text{ with respect to } w'\} - \frac{1}{2}(\#V(w) - \#V(w')).$$

This shows that the level of $G_{w \rightarrow w'}$ depends on the number of essential faces of w with respect to w' and therefore on the embedding of w' .

Proof The level of a red graph G is given by

$$I(G) = 2 \#V(G) - \#E(G) - \frac{1}{2} \sum_{f \in V(G)} \text{ed}(f).$$

By definition of $G_{w \rightarrow w'}$, we have

$$\{\text{essential faces of } w \text{ with respect to } w'\} = V(G_{w \rightarrow w'}).$$

The only thing to realize is that we have

$$2 \left(\#E(G)_{w \rightarrow w'} + \frac{1}{2} \sum_{f \in V(G_{w \rightarrow w'})} \text{ed}(f) \right) = \#V(w) - \#V(w'),$$

and this follows from the definition of $w_{G_{w \rightarrow w'}} = w'$. □

Definition 4.34 If f is a face of w , w' a simplification of w and r a region of w' , we say that f *avoids* r if $f \cap r = \emptyset$ or if the boundary of r in each connected component of $f \cap r$ joins two consecutive vertices of f (see Figure 37). In the first case we say that f *trivially avoids* r .

If f is an essential face of w with respect to w' and r is a region of w' , we say that f *fills* r if f does not avoid r . If F' is a set of regions of w' we say that f *fills* (resp. *avoids*) F' if it fills at least one region of F' (resp. avoids all the regions of F'). We define

$$n(f, F') \stackrel{\text{def}}{=} \#\{r \in F' \text{ such that } f \text{ fills } r\}.$$

If G' is a red graph of w' , we write $n(f, G')$ for $n(f, V(G'))$.




Figure 37: The local picture of a face f (in white) of w (in black) nontrivially avoiding a region r (in yellow) of w' (in orange).

With the same notation, and with a set of faces F of w ,

$$(2) \quad \#F = \#\{\text{faces } f \text{ of } F \text{ avoiding } F'\} + \sum_{f' \in F'} \sum_{\substack{f \in F \\ f \text{ fills } f'}} \frac{1}{n(f, F')}.$$

Lemma 4.35 *Let w be a ∂ -connected ε -web and w' a nice ∂ -connected simplification of w . Let F' be a collection of regions of w' . Then for every face f of w , $n(f, F') \leq 2$.*

Proof This is clear since $f \cap w'$ consists of at most one strand, so it intersects at most 2 regions of F' . □

Remark Let w be an ε -web, w' a nice ∂ -connected simplification of w and f an essential face of w with respect to w' . The simplification being nice, the face f intersects at most two regions of w' .

Let us consider a case where the face f intersects two regions r_1 and r_2 of w' . Either it (nontrivially) avoids one of them or fills both of them. Consider also a collection F' of regions of w' with $\{r_1, r_2\} \subseteq F'$. If f is a square then it (nontrivially) avoids one of the regions and $n(f, F') = 1$.

Suppose now that f has at least 6 edges. If f nontrivially avoids r_2 (and consequently $n(f, F') = 1$) then at least two neighbors (in $G_{w \rightarrow w'}$) of f fill r_1 (see Figure 38). If on the contrary f has just one neighbor which fills r_1 , then f fills r_2 . Under this condition, $n(f, F') = 2$.

Lemma 4.33 tells that in order to evaluate the index of a red graph coming from a simplification, it is crucial to control its number of vertices. The following definition gives a tool to count locally the number of vertices of a red graph.




Figure 38: On the left f avoids r_2 , on the right it fills r_1 and r_2 .

Definition 4.36 Let w an ε -web and w' a simplification of w , F a collection of faces of w , F' a set of regions of w' , and f' a face of w' . We set

$$\sigma(f', F \rightarrow F') \stackrel{\text{def}}{=} \sum_{\substack{f \in F \\ f \text{ fills } f'}} \frac{1}{n(f, F')} \in \frac{1}{2}\mathbb{Z}.$$

If G is a red graph for w and G' a red graph for w' , we write $\sigma(f', G \rightarrow G')$ for $\sigma(f', V(G) \rightarrow V(G'))$.

4F Proof of Lemma 4.22

In this section, we use the point of view developed in Section 4E to prove Lemma 4.22. We rewrite it with this new vocabulary.

Lemma 4.37 Let w be a ∂ -connected ε -web which contains no digon and exactly one square. We suppose furthermore that this square is adjacent to the unbounded region. Let G be a nice red graph of w and G' a nice red graph of $w' = w_G$. Then there exists a nice simplification \tilde{w} of w such that

- (A) the ε -webs $(w_G)_{G'}$ and \tilde{w} are isotopic,
- (B) $\#V(\tilde{G}) \geq \#V(G) + \#V(G')$, where \tilde{G} denotes the red graph $G_{w \rightarrow \tilde{w}}$.

Remark Note that the inequality (B) is what is needed, because, thanks to Lemma 4.33, this implies

$$I(\tilde{G}) \geq I(G) + I(G').$$

Proof Because of condition (A), the isotopy class of the web \tilde{w} is already known. To describe it completely, we only need to specify how \tilde{w} is embedded.

Let us denote by w'' the ε -web $(w')_{G'}$. For each face f' of w' which is a vertex of G' , let us denote by $N_{f'}$ a regular neighborhood of f' . In a regular neighborhood

of $w' \cap w''$, we define \tilde{w} to be equal to w'' . Therefore we only need to define the embedding of \tilde{w} in $\bigcup N_{f'}$.

If f' is a face of w' which is in $V(G')$, we distinguish two different cases:

- (i) The face f' corresponds to a vertex of G' with external degree equal to 0. The intersection of $N_{f'}$ with w'' is empty, and we define $\tilde{w} \cup N_{f'}$ to be the empty set. See [Figure 39](#) for an example.
- (ii) The face f' corresponds to a vertex of G' with external degree equal to 2. In this case, the intersection of $N_{f'}$ with w'' is a single strand. We define the intersection of \tilde{w} with $N_{f'}$ to be a single arc as well (with the same ends), but we still need to specify how this arc is embedded with respect to w . See [Figure 40](#) for an example.

These are the only cases to consider since G' is nice. It is clear that \tilde{w} and w'' are isotopic.

We claim that the following inequalities hold in case (i) and that it is possible to find embeddings of the arcs in case (ii) such that the inequalities hold for case (ii) as well:

$$(3) \quad \begin{cases} \sigma(f', \tilde{G} \rightarrow G') \geq \sigma(f', G \rightarrow G') + \frac{1}{2} & \text{if } S \subseteq N_{f'}, \\ \sigma(f', \tilde{G} \rightarrow G') \geq \sigma(f', G \rightarrow G') + 1 & \text{if } S \not\subseteq N_{f'}, \end{cases}$$

where S is the square of w .

This will be proven by [Lemma 4.39](#) and [Lemma 4.42](#).

The square S of w is in at most one $N_{f'}$, so if we sum (3) for all f' in $V(G')$, we obtain

$$\sum_{f' \in F'} \sigma(f', \tilde{G} \rightarrow G') \geq \sum_{f' \in F'} \sigma(f', G \rightarrow G') + \#V(G') - \frac{1}{2},$$

and using (2), we have

$$\#V(\tilde{G}) \geq \#V(G') + \#V(G) - \frac{1}{2},$$

because the set of faces of w which belong to $V(G)$ and avoid $V(G')$ is precisely the set of faces of w which belong to $V(\tilde{G})$ and avoid $V(G')$. But $\#V(\tilde{G})$ being an integer we have $\#V(\tilde{G}) \geq \#V(G') + \#V(G)$.

As we said, the inequalities (3) will be proven later, but let us examine the precise situation. The restriction of G to f' is a graph which satisfies the following conditions:

- It is bicolored (because the vertices of G are essential faces of w with respect to w').

- It is naturally embedded in a disk because $N_{f'}$ is diffeomorphic to a disk.
- The boundary of every face of G has at least 6 edges (this is a consequence of Proposition 4.15).
- The vertices in the interior of the disk have degree at least three (because the only possible square of w is adjacent to the unbounded region of w). If f is such a vertex of G , $n(f, V(G')) = 1$ (see the remark on page 1348).
- The vertices on the boundary (these are the ones which fill another region of w') have degree at least 1. If such a vertex f has degree 1 and is located next to another region of w' which is in $V(G')$, then $n(f, V(G')) = 2$ (see remark on page 1348) unless it is the square (in this case $n(f, V(G')) = 1$). In other cases $n(f, V(G')) \geq 1$. □




Figure 39: Example of the procedure to define \tilde{G} when the external degree of f' is equal to 0

4G Proof of combinatorial lemmas

This section is dedicated to the two technical lemmas used in the proof of Lemma 4.37. We first introduce certain ad-hoc objects and then state and prove the lemmas.

Definition 4.38 A D -graph is a graph G embedded into the disk D^2 . The set of vertices $V(G)$ is partitioned into two sets: $V^\partial(G)$ contains the vertices lying on ∂D^2 , while $V^{\text{in}}(G)$ contains the others. The set $F(G)$ of connected components of $D^2 \setminus G$ is partitioned into two sets: $F^{\text{in}}(G)$ contains the connected component included in $\overset{\circ}{D}^2$, while $F^\partial(G)$ contains the others.

A D -graph is said to be *nonelliptic* if

- every vertex v of $V^{\text{in}}(G)$ has degree greater than or equal to 3,
- every vertex v of $V^\partial(G)$ has degree greater than or equal than 1,
- the faces of $F^{\text{in}}(G)$ are of size⁸ at least 6.

A *colored D-graph* is a D -graph G together with

- a coloring (by **green** and **blue**) of the vertices of G such that two adjacent vertices have different colors (this implies that G is bipartite),
- a subdivision of ∂D^2 into two intervals (we allow one interval to be the empty set and the other one to be the full circle; in this case we say that G is *circle-colored*): a **green** one and a **blue** one (denoted by I_{blue} and I_{green}). When they are proper intervals we define x and y to be the two intersection points of I_{green} and I_{blue} with the convention that when we scan ∂D^2 clockwise, we see x , then I_{green} , then y and finally I_{blue} .

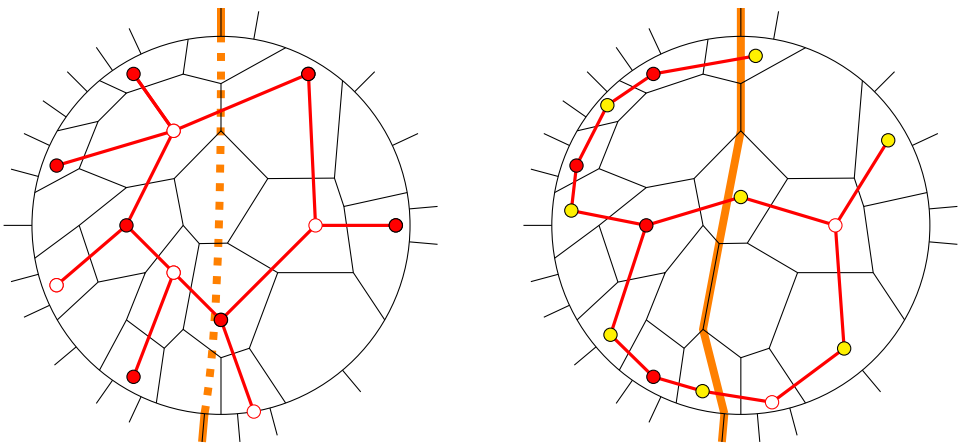


Figure 40: Example of the procedure to define \tilde{G} when the external degree of f' is equal to 2

The vertices of V^∂ are different from x and y . The color of a vertex does not need to match the color of the interval it lies on. We define V_{green} (resp. V_{blue}) to be the set of **green** (resp. **blue**) vertices and define $V_{\text{green}}^\partial$, $V_{\text{green}}^{\text{in}}$, V_{blue}^∂ and $V_{\text{blue}}^{\text{in}}$ similarly.

If G is a colored D -graph, and v is a vertex of V^∂ , we set

$$n(v) = \begin{cases} 2 & \text{if } v \text{ has degree 1 and the color of } v \text{ fits the color of the interval,} \\ 1 & \text{else.} \end{cases}$$

⁸We mean the number of edges of its boundary.

If v is a vertex of V^{in} , we set $n(v) = 1$. Note that this definition of n is a translation of the n of the previous section (see remark on page 1348). More precisely, for every vertex v but maybe the one corresponding to the square we have $n(v) \leq n_{\text{previous section}}(v)$.

4G.1 Case with external degree equal to 0

Lemma 4.39 *Let G be a nonelliptic circle-colored D -graph (with the circle colored by a color c). Then*

$$\#F \geq 1 + \sum_{v \in V_c} \frac{1}{n(v)}.$$

Remark This is exactly what we want because if we add $\sum_{v \notin V_c} 1/n(v)$ on both sides, we obtain the second inequality of (3).

Proof By symmetry, we may suppose that $c = \text{green}$. To prove the lemma, we consider the graph H obtained by gluing two copies of G along the boundary of D^2 . This is naturally embedded into the sphere. We write the Euler characteristic

$$(4) \quad \#F(H) - \#E(H) + \#V(H) = 1 + \#C(H),$$

where $C(H)$ is the set of connected components of H . We have the equalities

$$\begin{aligned} \#F(H) &= 2\#F^{\text{in}}(G) + \#F^\partial(G), \\ \#F^\partial(G) &= \#V^\partial(G) + 1 - \#C(H), \\ \#E(H) &= 2\#E(G) = \sum_{v \in V(G)} \deg(v) = 2 \sum_{v \in V_{\text{green}}(G)} \deg(v), \\ \#V(H) &= 2\#V^{\text{in}}(G) + \#V^\partial(G). \end{aligned}$$

Now we can rewrite (4) as

$$(5) \quad 2\#F^{\text{in}}(G) + 2\#F^\partial(G) + 2\#V^{\text{in}}(G) = 2 + 2\#E(G).$$

Now we use what we know about the degrees of the vertices:

$$\begin{aligned} \#E(G) &\geq \frac{3}{2}\#V^{\text{in}}(G) + \frac{1}{2}\#V^{\partial,1}(G) + \#V^{\partial,>1}(G), \\ \#E(G) &\geq 3\#V_{\text{green}}^{\text{in}}(G) + \#V_{\text{green}}^{\partial,1}(G) + 2\#V_{\text{green}}^{\partial,>1}(G). \end{aligned}$$

Here, $V^{\partial,1}$ (resp. $V^{\partial,>1}$) denotes the subset of V^∂ with degree equal to 1 (resp. with degree strictly greater than 1). If we sum $\frac{2}{3}$ of the first inequality and $\frac{1}{3}$ of the second




Figure 41: From left to right: a \cap , a λ and an H . The circle ∂D^2 is thick and gray, the D -graph is thin and black. Note that the vertices inside D^2 may have degree greater than 3.

one, and substitute this in (5), we obtain

$$\begin{aligned} \#F(G) + \#V^{\text{in}}(G) &\geq 1 + \#E(G), \\ \#F(G) + \#V^{\text{in}}(G) &\geq \#V^{\text{in}}(G) + \frac{1}{3}\#V^{\partial,1}(G) + \frac{2}{3}\#V^{\partial,>1}(G) \\ &\quad + \#V_{\text{green}}^{\text{in}}(G) + \frac{1}{3}\#V_{\text{green}}^{\partial,1}(G) + \frac{2}{3}\#V_{\text{green}}^{\partial,>1}(G), \\ \#F(G) &\geq \#V_{\text{green}}^{\text{in}}(G) + \frac{2}{3}\#V_{\text{green}}^{\partial,1}(G) + \frac{4}{3}\#V_{\text{green}}^{\partial,>1}(G) \\ &\quad + \frac{1}{3}\#V_{\text{blue}}^{\partial,1}(G) + \frac{2}{3}\#V_{\text{blue}}^{\partial,>1}(G) \\ &\geq \#V_{\text{green}}^{\text{in}}(G) + \frac{1}{2}\#V_{\text{green}}^{\partial,1}(G) + \#V_{\text{green}}^{\partial,>1}(G) \\ &\geq \sum_{v \in V_{\text{green}}} \frac{1}{n(v)}. \quad \square \end{aligned}$$

4G.2 Case with external degree equal to 2

Lemma 4.40 *If G is a nonelliptic D -graph, then all the faces of G are diffeomorphic to disks, and if it is nonempty, then at least one of the following situations happens:*

- (1) *The set $V^{\partial,>1}$ is nonempty.*
- (2) *There exist two \cap (see Figure 41) (if G consists of only one edge, then both \cap have the same edge, but are different, since one consider one side of the disk and the other).*
- (3) *There exist three λ or H (see Figure 41).*

Proof This is the same Euler characteristic argument that we used in Lemma 3.30. \square

Definition 4.41 A cut in a (non circle-) colored D -graph is a simple oriented path $\gamma : [0, 1] \rightarrow D$ such that

- $\gamma(0) = x$ and $\gamma(1) = y$ so that I_{green} is on the left and I_{blue} is on the right⁹ (see Figure 42);

⁹We use the convention that the left and right side are determined when one scans γ from x to y .

- for every face f of G , $f \cap \gamma$ is connected;
- the path γ crosses G either transversely at edges joining a **green** vertex on left and a **blue** vertex on the right, or at vertices of V^∂ whose colors do not match with the intervals they lie on.




Figure 42: A cut in a colored D -graph (note that G is elliptic)

If γ is a cut we denote by $l(\gamma)V(G)$ and $r(\gamma)V(G)$ the vertices located on the left (resp. right) of γ . (The vertices located on γ are meant to be both on the left and right).

Lemma 4.42 *Let G be a nonelliptic (non circle-) colored D -graph. Then there exists a cut γ such that*

$$\#F(G) \geq 1 + \sum_{v \in l(\gamma)V_{\text{green}}} \frac{1}{n(v)} + \sum_{v \in r(\gamma)V_{\text{blue}}} \frac{1}{n(v)}.$$

Remark This is exactly what we want because if we add

$$\sum_{v \in r(\gamma)V_{\text{green}}} \frac{1}{n(v)} + \sum_{v \in l(\gamma)V_{\text{blue}}} \frac{1}{n(v)}$$

on both sides, we obtain the second inequality of (3).

Proof The proof is done by induction on $s(G) \stackrel{\text{def}}{=} 3\#E(G) + 4\#V^{\partial, >1}(G)$. If this quantity is equal to zero then the D -graph is empty. In this case, we choose γ to be any simple arc joining x to y , and the lemma says $1 \geq 1$ which is true. We set

$$C(G, \gamma) \stackrel{\text{def}}{=} \sum_{v \in l(\gamma)V_{\text{green}}} \frac{1}{n(v)} + \sum_{v \in r(\gamma)V_{\text{blue}}} \frac{1}{n(v)}.$$




Figure 43: Local picture of G (left) and G' (right) in a regular neighborhood of v (or $v_1, \dots, v_{\deg(v)}$), when v is green and lies on I_{green}




Figure 44: Local picture of G (left) and G' (right) in a neighborhood of v (or v'), when v is blue and lies on I_{green}

It is enough to check the situations (1), (2) and (3) described in Lemma 4.40.

Situation (1) Let us denote by v a vertex of $V^{\partial, >1}$. There are two cases: the color of v fits with the colors of the intervals it lies on or not.

If the colors fit, say both are green, we consider G' the same colored D -graph as G but with v split into $v_1, v_2, \dots, v_{\deg(v)}$ all in $V^{\partial, 1}(G')$ (see Figure 43). We have $s(G') = S(G) - 4 < s(G)$ and G' nonelliptic, therefore we can apply the induction hypothesis. We can find a cut γ' with $\#F(G') \geq 1 + C(G', \gamma')$. Note that γ' does not cross any v' , so we can lift γ' to the D -graph G . This gives us γ . We have

$$\begin{aligned} C(G, \gamma) &= C(G', \gamma') + \frac{1}{n(v)} - \sum_{k=1}^{\deg(v)} \frac{1}{n(v_k)} \\ &= C(G', \gamma') + 1 - \frac{1}{2} \deg(v) \\ &\leq C(G', \gamma'). \end{aligned}$$

On the other hand $\#F(G) = \#F(G')$ so we have $\#F(G) \geq 1 + C(G, \gamma)$.

If the colors do not match (say v is blue), we construct a colored D -graph G' which is similar to G everywhere but in a regular neighborhood of v . The vertex v is pushed into the interior of D^2 (we denote it by v') and we add a new vertex v'' on ∂D^2 and an edge e joining v' and v'' . The colored D -graph G' is nonelliptic and $s(G') = s(G) - 4 + 3 < s(G)$ so we can apply the induction hypothesis and find a cut γ' with $\#F(G') \geq 1 + C(G', \gamma')$.



Figure 45: How to transform γ' into γ




Figure 46: The three possible configurations of G with two \cap

If γ' does not cross e we can lift γ' to G (this gives us γ). We have

$$C(G, \gamma) = C(G', \gamma') + \frac{1}{n(v)} - \frac{1}{n(v')} = C(G', \gamma') + 1 - 1 = C(G', \gamma').$$

On the other hand we have $F(G) = F(G')$, so $\#F(G) \geq 1 + C(G, \gamma)$.

Consider now the case where γ' crosses e . Then we consider the cut γ of G which is the same as γ' outside a regular neighborhood of v , and which around v crosses G in v (see Figure 45). We have

$$\begin{aligned} C(G, \gamma) &= C(G', \gamma') + \frac{1}{n(v)} - \frac{1}{n(v')} - \frac{1}{n(v'')} \\ &= C(G', \gamma') + 1 - 1 - \frac{1}{2} \\ &\leq C(G', \gamma'). \end{aligned}$$

But $\#F(G) = \#F(G')$, so we have $\#F(G) \geq 1 + C(G, \gamma)$.

Situation (2) We now suppose that G contains two \cap . Let us denote by v_g (resp. v_b) the green (resp. blue) vertex of the \cap and by e the edge of the cap. There are different possible configurations depending on the positions of x and y . As there are at least two caps, we may suppose y is outside a regular neighborhood of the \cap .

There are 3 different configurations (see Figure 46):

- The point x is far from the \cap .
- The point x is in the \cap and $v_g \in I_{\text{green}}$ and $v_b \in I_{\text{blue}}$.
- The point x is in the \cap and $v_g \in I_{\text{blue}}$ and $v_b \in I_{\text{green}}$.




Figure 47: How to transform γ' into γ




Figure 48: How to transform γ' into γ

We consider G' to be the colored D -graph similar to G except that the \cap is removed. The colored D -graph G' is nonelliptic and $s(G') = s(G) - 3 < s(G)$ so we can apply the induction hypothesis and find a cut γ' with $\#F(G') \geq 1 + C(G', \gamma')$.

Let us first suppose that x and the \cap are disjoint. Then v_b and v_g both lie either on I_{green} or on I_{blue} . By symmetry we may consider that they both lie on I_{green} . We can lift γ' to G (this gives γ) so that it does not meet the \cap . We have

$$C(G, \gamma) = C(G', \gamma') + \frac{1}{n(v_g)} = C(G', \gamma') + \frac{1}{2}.$$

But $\#F(G) = \#F(G') + 1$, hence $\#F(G) \geq 1 + C(G, \gamma)$.

Suppose now that the point x is in the \cap and $v_g \in I_{\text{green}}$ and $v_b \in I_{\text{blue}}$. We can lift γ' to G so that it crosses e (see Figure 47).

We have

$$C(G, \gamma) = C(G', \gamma') + \frac{1}{n(v_g)} + \frac{1}{n(v_b)} = C(G', \gamma') + \frac{1}{2} + \frac{1}{2} = C(G', \gamma') + 1.$$

But $\#F(G) = \#F(G') + 1$, hence $\#F(G) \geq 1 + C(G, \gamma)$.

Suppose now that the point x is in the \cap and $v_g \in I_{\text{blue}}$ and $v_b \in I_{\text{green}}$. We can lift γ' to G so that it crosses¹⁰ v_g (see Figure 48).

We have

$$C(G, \gamma) = C(G', \gamma') + \frac{1}{n(v_g)} = C(G', \gamma') + 1.$$

¹⁰We could have chosen to cross v_b .




Figure 49: On the top, the two possible configurations for a λ . On the bottom, the D -graphs G' obtained from G .

But $\#F(G) = \#F(G') + 1$, hence $\#F(G) \geq 1 + C(G, \gamma)$.

Situation (3) We suppose now that there are three λ or H . One can suppose that one of the λ 's or one of the H is disjoint from x and from y .

Let us first consider the case where there is a λ disjoint from x and y . Let us denote by v_1 and v_2 the two vertices of λ which belong to $V^\partial(G)$, by v the vertex of the λ which is in $V^{\text{in}}(G)$ and by e_1 (resp. e_2) the edge joining v and v_1 (resp. v_2). We consider G' to be the D -graph where the λ is replaced by a single strand: the edges e_1 and e_2 and the vertices v_1 and v_2 are suppressed. The vertex v is moved to ∂D^2 (and renamed v'). This is depicted in Figure 49. The colored D -graph G' is nonelliptic and $s(G') < s(G)$ so we can apply the induction hypothesis and find a cut γ' with $\#F(G') \geq 1 + C(G', \gamma')$.

The vertices v_1 and v_2 have the same color; by symmetry we may suppose that they are both green. This implies that v and v' are both blue.

There are two different configurations:

- The vertices v_1 and v_2 lie on I_{green} .
- The vertices v_1 and v_2 lie on I_{blue} .

Let us first suppose that the vertices v_1 and v_2 lie on I_{green} . If the cut γ' does not cross v' then we can canonically lift it to G . This gives us γ . We have

$$\begin{aligned} C(G, \gamma) &= C(G', \gamma') + \frac{1}{n(v_1)} + \frac{1}{n(v_2)} \\ &= C(G', \gamma') + \frac{1}{2} + \frac{1}{2}. \end{aligned}$$

But $\#F(G) = \#F(G') + 1$, hence $\#F(G) \geq 1 + C(G, \gamma)$.



Figure 50: How to transform γ' into γ



Figure 51: How to transform G into G'

If the cut γ' crosses v' , we lift γ' to G so that it crosses e_1 and e_2 (see Figure 48).

In this case,

$$\begin{aligned}
 C(G, \gamma) &= C(G', \gamma') + \frac{1}{n(v)} - \frac{1}{n(v')} + \frac{1}{n(v_1)} + \frac{1}{n(v_2)} \\
 &= C(G', \gamma') + 1 - 1 + \frac{1}{2} + \frac{1}{2}.
 \end{aligned}$$

Hence, $\#F(G) \geq 1 + C(G, \gamma)$.

Now suppose that the vertices v_1 and v_2 lie on I_{blue} . This implies that γ' does not meet v' , so we can lift γ' canonically to G . This gives us γ , and

$$\begin{aligned}
 C(G, \gamma) &= C(G', \gamma') + \frac{1}{n(v)} - \frac{1}{n(v')} \\
 &= C(G', \gamma') + 1 - 1.
 \end{aligned}$$

Hence $\#F(G) \geq 1 + C(G, \gamma)$.

We finally consider an H disjoint from x and y . We use the notation in Figure 51 to denote vertices and edges of H . We consider the D -graph G' obtained from G by simplifying H (see Figure 51 for details and notation). The colored D -graph G' is nonelliptic and $s(G') = s(G) - 3 \times 3 + 2 \times 4 < s(G)$ so we can apply the induction hypothesis and find a cut γ' with $\#F(G') \geq 1 + C(G', \gamma')$.

Up to symmetry there is only one configuration, therefore we may suppose that v_1 is green and lies on I_{green} . This implies that v_2 and v_3 are blue and that v_4 is green. Because of the color condition, the cut γ' does not cross v'_4 and may cross v'_3 . If it



Figure 52: How to transform γ' into γ

does not cross v'_3 , one can canonically lift γ' to G' , and

$$C(G, \gamma) = C(G', \gamma') + \frac{1}{n(v_1)} + \frac{1}{n(v_4)} - \frac{1}{n(v'_4)}$$

$$C(G, \gamma) \leq C(G', \gamma') + \frac{1}{2} + 1 - \frac{1}{2}.$$

$$C(G, \gamma) \leq C(G', \gamma') + 1.$$

But $\#F(G) = \#F(G') + 1$, hence $\#F(G) \geq 1 + C(G, \gamma)$.

If the cut γ' crosses v'_3 , one can lift it to G so that it crosses e_1 and e_2 (see Figure 52).

Hence

$$C(G, \gamma) = C(G', \gamma') + \frac{1}{n(v_1)} + \frac{1}{n(v_3)} - \frac{1}{n(v'_3)} + \frac{1}{n(v_4)} - \frac{1}{n(v'_4)}$$

$$C(G, \gamma) \leq C(G', \gamma') + \frac{1}{2} + 1 - 1 + 1 - \frac{1}{2}.$$

$$C(G, \gamma) \leq C(G', \gamma') + 1.$$

But $\#F(G) = \#F(G') + 1$, hence $\#F(G) \geq 1 + C(G, \gamma)$.

Conclusion In all situations, using the induction hypothesis we can construct a cut γ such that $\#F(G) \geq 1 + C(G, \gamma)$. This proves the lemma. \square

References

- [1] C Blanchet, N Habegger, G Masbaum, P Vogel, *Topological quantum field theories derived from the Kauffman bracket*, Topology 34 (1995) 883–927 [MR1362791](#)
- [2] P Freyd, D Yetter, J Hoste, W B R Lickorish, K Millett, A Ocneanu, *A new polynomial invariant of knots and links*, Bull. Amer. Math. Soc. 12 (1985) 239–246 [MR776477](#)
- [3] F Jaeger, *A new invariant of plane bipartite cubic graphs*, Discrete Math. 101 (1992) 149–164 [MR1172374](#)
- [4] M Khovanov, *A functor-valued invariant of tangles*, Algebr. Geom. Topol. 2 (2002) 665–741 [MR1928174](#)

- [5] **M Khovanov**, *sl(3) link homology*, *Algebr. Geom. Topol.* 4 (2004) 1045–1081
- [6] **M Khovanov, G Kuperberg**, *Correction to: “Web bases for $sl(3)$ are not dual canonical”*, *Pacific J. Math.* 188 (1999) 399 [MR1684195](#)
- [7] **J Kock**, *Frobenius algebras and 2D topological quantum field theories*, *LMS Student Texts* 59, Cambridge Univ. Press (2004) [MR2037238](#)
- [8] **G Kuperberg**, *Spiders for rank 2 Lie algebras*, *Comm. Math. Phys.* 180 (1996) 109–151 [MR1403861](#)
- [9] **A D Lauda, H Queffelec, E V Rose**, *Khovanov homology is a skew Howe 2–representation of categorified quantum $\mathfrak{sl}(m)$* [arXiv:1212.6076](#)
- [10] **M Mackaay, W Pan, D Tubbenhauer**, *The \mathfrak{sl}_3 –web algebra*, *Math. Z.* 277 (2014) 401–479 [MR3205780](#)
- [11] **M Mackaay, P Vaz**, *The universal \mathfrak{sl}_3 –link homology*, *Algebr. Geom. Topol.* 7 (2007) 1135–1169 [MR2336253](#)
- [12] **S Morrison, A Nieh**, *On Khovanov’s cobordism theory for \mathfrak{su}_3 knot homology*, *J. Knot Theory Ramifications* 17 (2008) 1121–1173 [MR2457839](#)
- [13] **J H Przytycki, P Traczyk**, *Invariants of links of Conway type*, *Kobe J. Math.* 4 (1988) 115–139 [MR945888](#)
- [14] **L-H Robert**, *A large family of indecomposable projective modules for the Khovanov–Kuperberg algebras of sl_3 –webs*, *J. Knot Theory Ramifications* 22 (2013) [MR3143579](#)
- [15] **L-H Robert**, *Sur l’homologie \mathfrak{sl}_3 des enchevêtrements; algèbre de Khovanov–Kuperberg*, PhD thesis, Université Paris 7 (2013) Available at <http://tinyurl.com/mhkv6bj>

MIN Fakultät, Fachbereich Mathematik
 Bundesstraße 55, 20146 Hamburg, Germany

louis-hadrien.robert@uni-hamburg.de

<http://www.math.uni-hamburg.de/home/robert/index.html>

Received: 13 September 2013 Revised: 8 August 2014