Journal of Combinatorial Theory, Series A 154 (2018) 129–144

Contents lists available at ScienceDirect

Journal of Combinatorial Theory, Series A

www.elsevier.com/locate/jcta

Torus link homology and the nabla operator

A.T. Wilson [1]

Department of Mathematics, University of Pennsylvania, Philadelphia, PA 19104, USA

ARTICLE INFO

Article history:
Received 19 September 2016
Available online xxxx

Keywords:
Torus links
Khovanov–Rozansky homology
Macdonald polynomials
Macdonald eigenoperators
Nabla operator

ABSTRACT

In recent work, Elias and Hogancamp develop a recurrence for the Poincaré series of the triply graded Khovanov–Rozansky homology of certain links, one of which is the (n, n) torus link. In this case, Elias and Hogancamp give a combinatorial formula for this homology that is reminiscent of the combinatorics of the modified Macdonald polynomial eigenoperator ∇. We give a combinatorial formula for the homologies of all complexes considered by Elias and Hogancamp. Our first formula is not easily computable, so we show how to transform it into a computable version. Finally, we conjecture a direct relationship between the (n, n) torus link case of our formula and the symmetric function ∇p_{1^n}.

Published by Elsevier Inc.

1. Introduction

We begin by establishing some notation from knot theory, following [6]. The remaining sections of the paper will take a more combinatorial perspective.

The *braid group on n strands*, denoted Br_n, can be defined by the presentation

$$\mathrm{Br}_n = \langle \sigma_1, \sigma_2, \ldots, \sigma_{n-1} \,|\, \sigma_i \sigma_{i+1} \sigma_i = \sigma_{i+1} \sigma_i \sigma_{i+1},\, \sigma_i \sigma_j = \sigma_j \sigma_i \rangle \tag{1}$$

E-mail address: andwils@math.upenn.edu.

[1] The author was supported by NSF Mathematical Sciences Postdoctoral Research Fellowship 1502512.

http://dx.doi.org/10.1016/j.jcta.2017.08.009
0097-3165/Published by Elsevier Inc.

for all $1 \leq i \leq n-2$ and $|i-j| \geq 2$. This group can be pictured as all ways to "braid" together n strands, where σ_i corresponds to crossing string $i+1$ over string i and the group operation is concatenation. One particularly notable braid is the *full twist braid* on n strands, denoted FT_n, which can be written

$$\mathrm{FT}_n = \left((\sigma_1)(\sigma_2\sigma_1)\dots(\sigma_{n-1}\sigma_{n-2}\dots\sigma_1)\right)^2 . \tag{2}$$

We will also need an operation ω on braids which corresponds to rotation around the horizontal axis. We define ω on Br_n by $\omega(\sigma_i) = \sigma_i$ and $\omega(\alpha\beta) = \omega(\beta)\omega(\alpha)$. Then ω is an anti-involution on Br_n. A braid that has the property that the string that begins in column i also ends in column i for all i is called a *pure braid*.

Given a braid with n strands, one can form a *link* (i.e. nonintersecting collection of knots) by identifying the top of the strand that begins in position i with the bottom of the strand that ends in position i for $1 \leq i \leq n$. The result is called a *closed braid*. Alexander proved that every link can be represented by a closed braid (although this representation is not unique) [1]. The closure of a pure braid is a link that consists of n separate unknots linked together.

In [6], Elias and Hogancamp assign a diagram to every binary word v. We describe this assignment here – see Fig. 1 for an example. Say $v \in \{0,1\}^n$ with $|v| = m$. We begin with n strands at the top of the diagram and two (currently unlabeled) boxes at the bottom of the diagram. The left box has $n-m$ inputs and outputs and the right box has m inputs and outputs. For $i = 1$ to n, if $v_i = 1$ we feed string i into the leftmost available input in the right box; otherwise, we feed string i into the leftmost available input in the left box. All crossings that occur are forced to be "positive," i.e. the right strand crosses over the left strand. We call the braid that is formed by feeding the strands into the boxes in this manner β_v. We prepend $\omega(\beta_v)$ to β_v to obtain our final diagram.

Next, we transform the diagram to a complex C_v of Soergel bimodules by replacing the left box with the complex associated to the full twist braid FT_{n-m} and the right box with the symmetrizer K_m [12]. We note that C_{0^n} is the bimodule associated to FT_n and that the closure of that associated braid is the (n,n) torus link. The combinatorics of other links, in particular the (m,n) torus link for m and n coprime, has been studied by a variety of authors in recent years [8,7]. Haglund gives an overview of this work from a combinatorial perspective in [10].

Elias and Hogancamp consider the *Hochschild homology* of this C_v complex of bimodules; this is sometimes called Khovanov–Rozansky homology [13,14]. This homology has three gradings: the bimodule degree (using the variable Q), the homological degree (T), and the Hochschild degree (A). After the grading shifts $q = Q^2$, $t = T^2Q^{-2}$, and $a = AQ^{-2}$, Elias and Hogancamp give a recurrence for the Poincaré series of this triply graded homology, which they denote $f_v(q,t,a)$. We will use their recurrence to define $f_v(q,t,a)$ in Section 2. They also give a combinatorial formula for the special case $f_{0^n}(q,t,a)$. We will give two combinatorial formulas for $f_v(q,t,a)$ for every $v \in \{0,1\}^n$.

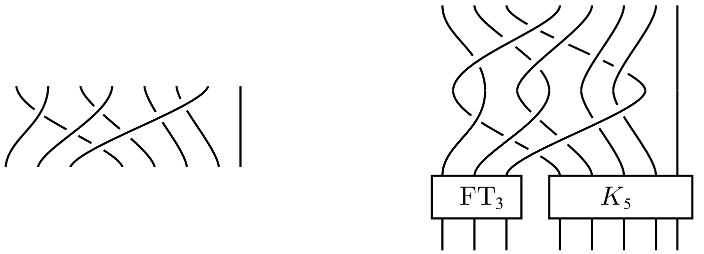

Fig. 1. For $v = 10101101$ we have drawn the braid β_v on the left and the complex C_v on the right. FT$_3$ is the complex of the full twist braid and K_5 is a certain complex defined recursively in [6]. Images are used courtesy of [6].

In Section 2, we define a symmetric function $L_v(x; q, t)$ which we call the *link symmetric function*. Its definition is reminiscent of the combinatorics of the Macdonald eigenoperator ∇, introduced in [3]. We prove that $f_v(q, t, a)$ is equal to a certain inner product with $L_v(x; q, t)$.

The main weakness of our first formula is that it is a sum over infinitely many objects, so it is not clear how to compute using this formula. We address this issue in Section 3, obtaining a finite formula for $L_v(x; q, t)$ using a collection of combinatorial objects we call *barred Fubini words*.

We close by presenting some conjectures in Section 4. In particular, we conjecture that

$$L_{0^n}(x; q, t) = (1 - q)^{-n} \nabla p_{1^n}, \tag{3}$$

where the terminology is defined in Section 4. A proof of this conjecture would provide the first combinatorial interpretation for ∇p_{1^n}. There has been much recent work establishing combinatorial interpretations for ∇e_n [5] and ∇p_n [18]. We believe that $L_v(x; q, t)$ is also related to Macdonald polynomials for general v, although we do not have an explicit conjecture in this direction.

2. An infinite formula

Let $\mathbb{N} = \{0, 1, 2, \ldots\}$ and $\mathbb{P} = \{1, 2, 3, \ldots\}$. We begin by defining two statistics.

Definition 2.1. Given words $\gamma \in \mathbb{N}^n$ and $\pi \in \mathbb{P}^n$, we define

$$\operatorname{area}(\gamma) = |\gamma| - \#\{1 \leq i \leq n : \gamma_i > 0\} \tag{4}$$

$$\begin{aligned} \operatorname{dinv}(\gamma, \pi) = \#\{1 \leq i < j \leq n : \gamma_i = \gamma_j, \pi_i > \pi_j\} \\ + \#\{1 \leq i < j \leq n : \gamma_i + 1 = \gamma_j, \pi_i < \pi_j\} \end{aligned} \tag{5}$$

$$x^\pi = \prod_{i=1}^{n} x_{\pi_i}. \tag{6}$$

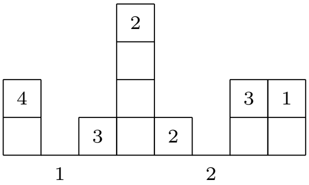

Fig. 2. We have depicted the example $\gamma = 20141022$ and $\pi = 41322231$ by drawing bottom-justified columns with heights $\gamma_1, \gamma_2, \ldots, \gamma_8$ and the labels π_i are placed as high as possible in each column. In this example, we compute $\operatorname{area}(\gamma) = 6$, $\operatorname{dinv}(\gamma, \pi) = 7$, where the contributing pairs are in columns $(1, 7)$, $(1, 8)$, $(2, 3)$, $(2, 5)$, $(3, 5)$, $(5, 7)$, $(7, 8)$, and $x^\pi = x_1^2 x_2^3 x_3^2 x_4$.

In Fig. 2, we draw a diagram for $\gamma = 20141022$ and $\pi = 41322231$. Area counts the empty boxes in such a diagram, dinv counts certain pairs of labels, and x^π records all labels that appear in the diagram.

Definition 2.2. Given $n \in \mathbb{P}$ and $v \in \{0,1\}^n$, define

$$L_v = L_v(x; q, t) = \sum_{\substack{\gamma \in \mathbb{N}^n,\, \pi \in \mathbb{P}^n \\ \gamma_i = 0 \Leftrightarrow v_i = 1}} q^{\operatorname{area}(\gamma)} t^{\operatorname{dinv}(\gamma,\pi)} x^\pi. \tag{7}$$

Perhaps the first thing to note about L_v is that it can be expressed as a sum of LLT polynomials [15]; as a result, it is symmetric in the x_i variables. More precisely, each $\gamma \in \mathbb{N}^n$ can be associated with an n-tuple $\lambda(\gamma)$ of single cell partitions in the plane, where the ith cell is placed on diagonal γ_i and the order is not changed. Using the notation of [11], the unicellular LLT polynomial $G_{\lambda(\gamma)}(x; t)$ can be used to write

$$L_v = \sum_{\substack{\gamma \in \mathbb{N}^n \\ \gamma_i = 0 \Leftrightarrow v_i = 1}} q^{\operatorname{area}(\gamma)} G_{\lambda(\gamma)}(x; t). \tag{8}$$

Since LLT polynomials are symmetric, every L_v is also symmetric.

We also remark that L_{1^n} is equal to the modified Macdonald polynomial $\widetilde{H}_{1^n}(x; q, t)$, which is also equal to the graded Frobenius series of the coinvariants of $\mathfrak{S}_n$ with grading in t.

Next, we note that the Poincaré series $f_v(q, t, a)$ can be recovered as a certain inner product of L_v. We follow the standard notation for symmetric functions and their usual inner product, as described in Chapter 7 of [20]. Before we can prove Theorem 2.1, we need the following lemma.

Lemma 2.1.

$$L_{0^n} = \frac{1}{1-q} L_{10^{n-1}}. \tag{9}$$

Proof. By definition,

$$L_{0^n} = \sum_{\gamma,\pi \in \mathbb{P}^n} q^{\mathrm{area}(\gamma)} t^{\mathrm{dinv}(\gamma,\pi)} x^{\pi}. \tag{10}$$

Our aim is to show that

$$L_{0^n} = q^n L_{0^n} + \left(1 + q + \ldots + q^{n-1}\right) L_{10^{n-1}} \tag{11}$$

which clearly implies the lemma.

If $\gamma_i > 1$ for all i, then let γ' be the word obtained by decrementing each entry in γ by 1. Set $\pi' = \pi$. Note that the pair (γ', π') has

$$\mathrm{area}(\gamma') = \mathrm{area}(\gamma) - n \tag{12}$$

$$\mathrm{dinv}(\gamma', \pi') = \mathrm{dinv}(\gamma, \pi) \tag{13}$$

$$x^{\pi'} = x^{\pi}. \tag{14}$$

Furthermore, every pair of words of positive integers can be obtained as (γ', π') in this fashion. This case corresponds to the first term on the right-hand side of (11).

The other case we must consider is if $\gamma_i = 1$ for some i. Let k be the rightmost position such that $\gamma_k = 1$. Then we define

$$\gamma'' = (\gamma_k - 1)(\gamma_{k+1} - 1) \ldots (\gamma_n - 1)\gamma_1\gamma_2 \ldots \gamma_{k-1} \tag{15}$$

$$\pi'' = \pi_k \pi_{k+1} \ldots \pi_n \pi_1 \pi_2 \ldots \pi_{k-1}. \tag{16}$$

It is straightforward to check that

$$\mathrm{area}(\gamma'') = \mathrm{area}(\gamma) - (n - k) \tag{17}$$

$$\mathrm{dinv}(\gamma'', \pi'') = \mathrm{dinv}(\gamma, \pi) \tag{18}$$

$$x^{\pi''} = x^{\pi}. \tag{19}$$

Furthermore, by construction we have $\gamma''_1 = 0$ and the other entries of γ'' are greater than 0. Summing over all values of k and pairs (γ'', π'') obtained in this way, we get the remaining terms in the right-hand side of (11). □

In [6], the authors prove that $f_v(q,t,a)$ satisfies a certain recurrence. We will use their recurrence as our definition of $f_v(q,t,a)$.

Definition 2.3. Given $v \in \{0,1\}^n$ and $w \in \{0,1\}^{n-|v|}$, we form a word $u \in \{0,1,2\}^n$ that depends on v and w. We set $u_i = 1$ if $v_i = 1$. If $v_i = 0$, say that we are at the jth zero in v, counting from left to right. Then we set $u_i = 2w_j$. For example, if $v = 10110100$ and $w = 0110$ then $u = 10112120$. We form a product

$$P_{v,w}(t,a) = \prod_{i\,:\,v_i=1} \left(t^{\#\{j<i\,:\,u_j=1\}+\#\{j>i\,:\,u_j=2\}} + a \right). \tag{20}$$

Using the recurrence in [6], we define

$$f_v(q,t,a) = \sum_{w\in\{0,1\}^{n-|v|}} q^{n-|v|-|w|} P_{v,w}(t,a) f_w(q,t,a) \tag{21}$$

with base cases $f_\emptyset(q,t,a) = 1$ and $f_{0^n}(q,t,a) = (1-q)^{-1} f_{10^{n-1}}(q,t,a)$.

Theorem 2.1. *For any $v \in \{0,1\}^n$,*

$$f_v(q,t,a) = \sum_{d=0}^{n} \langle L_v, e_{n-d}h_d \rangle\, a^d. \tag{22}$$

Proof. Let us denote the right-hand side of the statement in the theorem by $\ell_v(q,t,a)$. The goal of this proof is to show that $\ell_v(q,t,a)$ satisfies (21). As discussed in Chapter 6 of [9], taking the inner product with $e_{n-d}h_d$ can be thought of as replacing the word $\pi \in \mathbb{P}^n$ in Definition 2.2 with a word π containing $n-d$ $\underline{0}$'s and d 1's. For the purposes of computing $\mathrm{dinv}(\gamma,\pi)$ we consider $\underline{0}$ to be less than itself, but we do not make this convention for 1. For example, if $\gamma = 1111$ and $\pi = \underline{0}1\underline{0}1$, we have $\mathrm{dinv}(\gamma,\pi) = 2$, where the two pairs we count are $(1,3)$ and $(1,2)$. With these definitions, we can write

$$\ell_v(q,t,a) = \sum_{\substack{\gamma\in\mathbb{N}^n,\, \pi\in\{\underline{0},1\}^n \\ \gamma_i=0 \Leftrightarrow v_i=1}} q^{\mathrm{area}(\gamma)} t^{\mathrm{dinv}(\gamma,\pi)} a^{\#1\text{'s in } \pi}. \tag{23}$$

Now we proceed by induction. We can simply define $\ell_\emptyset(q,t,a) = 1$ and the fact that $\ell_{0^n}(q,t,a) = (1-q)^{-1}\ell_{10^{n-1}}(q,t,a)$ follows from Lemma 2.1.

Given a word $\gamma \in \mathbb{N}^n$, we form a word u by setting $u_i = 1$ if $\gamma_i = 0$, $u_i = 2$ if $\gamma_i = 1$, and $u_i = 0$ otherwise. From this word u we construct another word $w \in \{0,1\}^{n-|v|}$ by scanning u from left to right and appending a 1 to w whenever we see a 2 in u and appending a 0 to w whenever we see a 0 in u. For example, if $\gamma = 013021$ we have $u = 120102$ and $w = 1001$.

Now we can explain why $\ell_v(q,t,a)$ satisfies (21). First, we note that the $q^{n-|v|-|w|}$ term counts the contribution of empty boxes in row 1 to area. We also claim that $P_{v,w}(t,a)$ uniquely counts the contributions from dinv pairs (i,j) with either $\gamma_i = \gamma_j = 0$ or $\gamma_i = 0$ and $\gamma_j = 1$. For each such pair, say that the pair *projects onto* j if $\gamma_i = \gamma_j = 0$ or i if $\gamma_i = 0$ and $\gamma_j = 1$. Then every such pair projects onto a unique i such that $\gamma_i = 0$, which is equivalent to $v_i = 1$. Furthermore, the number of pairs projecting onto a particular i is 0 if $\pi_i = 1$ and

$$\#\{j<i\,:\,\gamma_j=0\} + \#\{j>i\,:\,\gamma_j=1\} = \#\{j<i\,:\,u_j=1\} + \#\{j>i\,:\,u_j=2\} \tag{24}$$

if $\pi_i = \underline{0}$. Hence, $P_{v,w}(t,a)$ accounts for the contribution all such dinv pairs. By induction, $\ell_w(q,t,a)$ accounts for all other area and all other dinv pairs.

For example, if $v = 100100$ and $\gamma = 013021$ then we have the (currently unlabeled) diagram

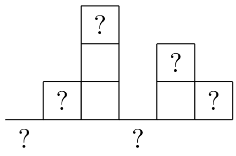

where the question marks will be replaced by labels $\overline{0}$ and 1. We noted above that $u = 120102$ and $w = 1001$. This example contributes to the term

$$q^{6-2-2}P_{100100,1001}(t,a)\ell_{1001}(q,t,a) = q^2(t^2+a)^2\ell_{1001}(q,t,a). \tag{25}$$

in the recurrence for $\ell_v(q,t,a)$. The q^2 term comes from the lowest boxes in columns 3 and 5, since we know that these boxes will not contain any labels (and will therefore contribute powers of q). The t^2+a factors correspond to placing a $\overline{0}$ or 1 in columns 1 and 4. If we place a $\overline{0}$ in column 1, then we will have dinv pairs $(1,2)$ and $(1,6)$ that both project onto 1. If we place a 1 in column 1, no dinv pairs will project onto column 1 but we will get another power of a. Similarly, if column 4 contains a $\overline{0}$ then we get dinv pairs $(1,4)$ and $(4,6)$ which project onto 4; otherwise, no dinv pairs project onto column 4. Finally, the $\ell_{1001}(q,t,a)$ accounts for everything that occurs in higher rows of the diagram.

The $v = 0^n$ case of the result follows from Lemma 2.1. □

For the sake of comparison with [6], we give a simplified formula that directly computes $f_v(q,t,a)$ from Theorem 2.1. Given $\gamma \in \mathbb{N}^n$ and $1 \le i \le n$, let

$$\mathrm{dinv}_i(\gamma) = \#\{j < i : \gamma_j = \gamma_i\} + \#\{j > i : \gamma_j = \gamma_i + 1\}. \tag{26}$$

Corollary 2.1.

$$f_v(q,t,a) = \sum_{\substack{\gamma\in\mathbb{N}^n \\ \gamma_i=0 \Leftrightarrow v_i=1}} q^{\mathrm{area}(\gamma)} \prod_{i=1}^{n} \left(a + t^{\mathrm{dinv}_i(\gamma)}\right) \tag{27}$$

where, as before, $\mathrm{area}(\gamma) = |\gamma| - \#\{1 \le i \le n : \gamma_i > 0\}$.

If $v = 0^n$ and $a = 0$, this is exactly Theorem 1.9 in [6].

3. A finite formula

Although the combinatorial definition of L_v is straightforward, it is not computationally effective[2] since it is a sum over infinitely many words $\gamma \in \mathbb{N}^n$. We rectify this issue in Theorem 3.1 below. The idea is to compress the vectors γ while altering the statistics so that the link polynomial L_v is not changed.

Definition 3.1. A word $\gamma \in \mathbb{N}^n$ is a *Fubini word* if every integer $0 \leq k \leq \max(\gamma)$ appears in γ.

For example, 41255103 is a Fubini word but 20141022 is not a Fubini word, since it contains a 4 but not a 3. We call these Fubini words because they are counted by the Fubini numbers ([19], A000670), which also count ordered partitions of the set $\{1, 2, \ldots, n\}$. We will actually be interested in certain decorated Fubini words.

Definition 3.2. Given $v \in \{0,1\}^n$, we say that a Fubini word γ is *associated* with v if either

- $v = 0^n$ and the only zero in γ occurs at γ_1, or
- $v \neq 0^n$ and $\gamma_i = 0$ if and only if $v_i = 1$.

Definition 3.3. A *barred Fubini word* associated with v is a Fubini word γ associated with v where we may place bars over certain entries. Specifically, the entry γ_j may be barred if

(1) $\gamma_j > 0$,
(2) γ_j is unique in γ, and
(3) for each $i < j$ we have $\gamma_i < \gamma_j$, i.e. γ_j is a left-to-right maximum in γ.

We denote the collection of barred Fubini words associated with v by $\overline{\mathcal{F}}_v$.

For example,

$$\overline{\mathcal{F}}_0 = \{0\} \tag{28}$$

$$\overline{\mathcal{F}}_{00} = \{01, 0\overline{1}\} \tag{29}$$

$$\overline{\mathcal{F}}_{000} = \{011, 012, 0\overline{1}2, 01\overline{2}, 0\overline{1}\overline{2}, 021, 0\overline{2}1\}. \tag{30}$$

The sequence $|\overline{\mathcal{F}}_{0^n}|$ for $n \in \mathbb{N}$ begins $1, 1, 2, 7, 35, 226, \ldots$ and seems to appear in the OEIS as A014307 [19]. One way to define sequence A014307 is that it has exponential generating function

[2] There are also infinitely many $\pi \in \mathbb{P}^n$, but this problem can be rectified with standardization [9].

v	$\overline{\mathcal{F}}_v$
111	000
011	$100, \bar{1}00$
101	$010, 0\bar{1}0$
110	$001, 00\bar{1}$
001	$110, 120, 1\bar{2}0, \bar{1}20, \bar{1}\bar{2}0, 210, \bar{2}10$
010	$101, 102, 10\bar{2}, \bar{1}02, \bar{1}0\bar{2}, 201, \bar{2}01$
100	$011, 012, 0\bar{1}2, 01\bar{2}, 0\bar{1}\bar{2}, 021, 0\bar{2}1$
000	$011, 012, 0\bar{1}2, 01\bar{2}, 0\bar{1}\bar{2}, 021, 0\bar{2}1$

Fig. 3. We have listed the barred Fubini words $\overline{\mathcal{F}}_v$ for each $v \in \{0,1\}^3$.

$$\sqrt{\frac{e^z}{2-e^z}}. \tag{31}$$

This sequence is given several combinatorial interpretations in [17]. It would be interesting to obtain a bijection between $\overline{\mathcal{F}}_{0^n}$ and one of the collections of objects in [17]. See Fig. 3 for more examples of barred Fubini words.

Given a barred Fubini word γ and a word $\pi \in \mathbb{P}^n$, we modify the dinv statistic slightly:

$$\begin{aligned} \operatorname{dinv}(\gamma,\pi) = \#\{1 \leq i < j \leq n : \gamma_i = \gamma_j, \pi_i > \pi_j\} \\ + \#\{1 \leq i < j \leq n : \gamma_i + 1 = \gamma_j, \pi_i < \pi_j, \gamma_j \text{ is not barred}\}. \end{aligned} \tag{32}$$

We also let $\operatorname{bar}(\gamma)$ be the number of barred entries in γ. We have the following result.

Theorem 3.1. *For $v \in \{0,1\}^n$,*

$$L_v = \sum_{\substack{\gamma \in \overline{\mathcal{F}}_v \\ \pi \in \mathbb{P}^n}} q^{\operatorname{area}(\gamma)+\operatorname{bar}(\gamma)} t^{\operatorname{dinv}(\gamma,\pi)} (1-q)^{-\operatorname{bar}(\gamma)-\chi(v=0^n)} x^\pi \tag{33}$$

where χ of a statement is 1 if the statement is true and 0 if it is false.

Proof. Assume, for now, that $v \neq 0^n$. Let $\overline{\mathcal{F}}_v^{(0)}$ denote the set of all $\gamma \in \mathbb{N}^n$ such that $\gamma_i = 0$ if and only if $v_i = 1$. These are exactly the γ that appear in the definition of L_v, Definition 2.2. Note that $\overline{\mathcal{F}}_v^{(0)}$ is *not* a subset of $\overline{\mathcal{F}}_v$; rather, $\overline{\mathcal{F}}_v \subseteq \overline{\mathcal{F}}_v^{(0)}$. For each $1 \leq k \leq n$, let $\overline{\mathcal{F}}_v^{(k)}$ be the set of vectors $\gamma \in \mathbb{N}^n$ decorated with bars such that

(1) $\gamma_i = 0$ if and only if $v_i = 1$,
(2) each $i \in \{0, 1, \ldots, \max(\gamma)\}$ such that $i \leq k$ appears in γ,

where $\gamma_j \in \overline{\mathcal{F}}_v^{(k)}$ may be barred if

(1) $0 < \gamma_j \leq k$,
(2) γ_j is unique in γ, and
(3) for each $i < j$ we have $\gamma_i < \gamma_j$, i.e. γ_j is a left-to-right maximum in γ.

Note that $\overline{\mathcal{F}}_v^{(0)} \supseteq \overline{\mathcal{F}}_v^{(1)} \supseteq \ldots \supseteq \overline{\mathcal{F}}_v^{(n)} = \overline{\mathcal{F}}_v$; in particular, this means $\overline{\mathcal{F}}_v^{(n)}$ is finite. The goal of the proof is to interpolate from $\overline{\mathcal{F}}_v^{(0)}$ to $\overline{\mathcal{F}}_v^{(n)} = \overline{\mathcal{F}}_v$. For convenience, we set

$$\mathrm{wt}_{\gamma,\pi} = \mathrm{wt}_{\gamma,\pi}(x;q,t) = q^{\mathrm{area}(\gamma)+\mathrm{bar}(\gamma)} t^{\mathrm{dinv}(\gamma,\pi)} (1-q)^{-\mathrm{bar}(\gamma)} x^{\pi}, \tag{34}$$

where the dinv statistic is the one we defined for barred Fubini words. Our goal is to show that

$$\sum_{\substack{\gamma \in \overline{\mathcal{F}}_v^{(k-1)} \\ \pi \in \mathbb{P}^n}} \mathrm{wt}_{\gamma,\pi} = \sum_{\substack{\gamma \in \overline{\mathcal{F}}_v^{(k)} \\ \pi \in \mathbb{P}^n}} \mathrm{wt}_{\gamma,\pi} \tag{35}$$

for each $1 \leq k \leq n$. Then we can chain together these identities for $k = 1, 2, \ldots, n$ to obtain the desired result.

First, we remove the intersection $\overline{\mathcal{F}}_v^{(k-1)} \cap \overline{\mathcal{F}}_v^{(k)}$ from both summands in (35) to obtain the equivalent statement

$$\sum_{\substack{\gamma \in \overline{\mathcal{F}}_v^{(k-1)} \setminus \overline{\mathcal{F}}_v^{(k)} \\ \pi \in \mathbb{P}^n}} \mathrm{wt}_{\gamma,\pi} = \sum_{\substack{\gamma \in \overline{\mathcal{F}}_v^{(k)} \setminus \overline{\mathcal{F}}_v^{(k-1)} \\ \pi \in \mathbb{P}^n}} \mathrm{wt}_{\gamma,\pi} \,. \tag{36}$$

Now we wish to describe the γ that appear in the left- and right-hand summands of (36). $\gamma \in \overline{\mathcal{F}}_v^{(k-1)}$ is not in $\overline{\mathcal{F}}_v^{(k)}$ if and only if $k < \max(\gamma)$ and γ does not contain a k. On the other hand, $\gamma \in \overline{\mathcal{F}}_v^{(k)}$ is not in $\overline{\mathcal{F}}_v^{(k-1)}$ if and only if it contains a single k and that k is barred. This allows us to rewrite (36) as

$$\sum_{\substack{\gamma \in \overline{\mathcal{F}}_v^{(k-1)} \\ k \notin \gamma \\ k < \max(\gamma) \\ \pi \in \mathbb{P}^n}} \mathrm{wt}_{\gamma,\pi} = \sum_{\substack{\gamma \in \overline{\mathcal{F}}_v^{(k)} \\ \overline{k} \in \gamma \\ \pi \in \mathbb{P}^n}} \mathrm{wt}_{\gamma,\pi} \,. \tag{37}$$

Specifically, for each subset $S \subseteq \{1, 2, \ldots, n\}$ we will show that

$$\sum_{\substack{\gamma \in \overline{\mathcal{F}}_v^{(k-1)} \\ k \notin \gamma \\ k < \max(\gamma) \\ \gamma_i < k \Leftrightarrow i \in S \\ \pi \in \mathbb{P}^n}} \mathrm{wt}_{\gamma,\pi} = \sum_{\substack{\gamma \in \overline{\mathcal{F}}_v^{(k)} \\ \overline{k} \in \gamma \\ \gamma_i < k \Leftrightarrow i \in S \\ \pi \in \mathbb{P}^n}} \mathrm{wt}_{\gamma,\pi} \,. \tag{38}$$

Then summing over all S will conclude the proof.

We consider the left-hand side of (38). Note that there cannot be any dinv between columns i and j if $\gamma_i < k$ and $\gamma_j > k$. In this sense, the columns i with $\gamma_i < k$ are independent of the columns j with $\gamma_j > k$. This allows us to write the left-hand side of (38) as a product

$$q^{n-|S|} L_{0^{n-|S|}} F_{v,S} \tag{39}$$

where $F_{v,S}$ is a certain symmetric function that accounts for all contribution to the weights coming from columns $i \in S$. Since $k < \max(\gamma)$, we know that $|S| < n$. The factor of q appears because each of the entries $j \notin S$ has an empty box in the diagram that is not counted by either of the other factors. Now we can use Lemma 2.1 to rewrite this product as

$$\frac{q^{n-|S|}}{1-q} L_{10^{n-|S|-1}} F_{v,S}. \tag{40}$$

We switch our attention to the right-hand side of (38), which we would like to show is equal to (40). Let m be minimal such that $m \notin S$ in the right-hand side of (38), i.e. $\gamma_m \geq k$. Since $\overline{k} \in \gamma$, such an m must exist. Furthermore, we must have $\gamma_m = \overline{k}$. We know that $\overline{k} \in \gamma$, and the furthest left it can appear in γ is in column m. If $\gamma_m = k$ without a bar or $\gamma_m > k$ then $\overline{k}$ cannot appear to the right of column m by conditions (2) and (3) for the barring of entries in $\overline{\mathcal{F}}_v^{(k)}$.

We note that, by the definition of dinv for barred words, there are no dinv pairs (i, j) with $i \in S$ and $j \notin S$, i.e. $\gamma_i < k$ and $\gamma_j \geq k$ for γ that appear in the sum on the right-hand side of (38). Hence, the contributions from columns in S will be independent of the contributions from columns not in S. The columns in S contribute $F_{v,S}$. The columns $j \notin S$ give

$$q^{n-|S|-1} \frac{q}{1-q} L_{10^{n-|S|-1}} = \frac{q^{n-|S|}}{1-q} L_{10^{n-|S|-1}}. \tag{41}$$

To see this, note that we have one fewer extra box than in the previous case; specifically, we have lost the extra box in column m, yielding a factor of $q^{n-|S|-1}$. We have a new bar in column m which contributes a factor of $q/(1-q)$. Multiplying these terms with $F_{v,S}$, we obtain (40).

Finally, we must address the case $v = 0^n$. In this case, we immediately use $L_{0^n} = (1-q)^{-1} L_{10^{n-1}}$ and then proceed as above. This is why Fubini words associated with 0^n have an "extra" zero at the beginning. This also slightly adjusts the weight of the summands, explaining the $\chi(v = 0^n)$ part of the theorem. □

As in Section 2, we give a formula for computing $f_v(q, t, a)$ directly. Given a barred Fubini word γ, we define

$$\mathrm{dinv}_i(\gamma) = \#\{j < i : \gamma_j = \gamma_i\} + \#\{j > i : \gamma_j = \gamma_i + 1, \gamma_j \text{ is not barred}\}. \tag{42}$$

Corollary 3.1.

$$f_v(q, t, a) = \sum_{\gamma \in \overline{\mathcal{F}}_v} q^{\mathrm{area}(\gamma) + \mathrm{bar}(\gamma)} (1-q)^{-\mathrm{bar}(\gamma) - \chi(v = 0^n)} \prod_{i=1}^{n} \left(a + t^{\mathrm{dinv}_i(\gamma)}\right). \tag{43}$$

t^2			
t	qt	q^2t	
1	q	q^2	q^3

Fig. 4. This is the Ferrers diagram of the partition $\mu = (4, 3, 1)$. In each cell we have written the monomial $q^i t^j$ that corresponds to the cell, yielding $B_\mu = \{1, q, q^2, q^3, t, qt, q^2t, t^2\}$.

4. Conjectures

So far, we have used the inner product $\langle L_v, e_{n-d}h_d \rangle$ to compute $f_v(q, t, a)$; one might wonder if there is any value in studying the full symmetric function L_v. In this section, we conjecture that the link symmetric function L_v is closely related to the combinatorics of Macdonald polynomials, hinting at a stronger connection between Macdonald polynomials and link homology. Following [6], we must first define a "normalized" version of the link symmetric function L_v.

Definition 4.1.

$$\widetilde{L}_v = \widetilde{L}_v(x; q, t) = (1 - q)^{n-|v|} L_v(x; q, t). \tag{44}$$

We could also define $\widetilde{L}_v$ in terms of diagrams; each box that contains a number contributes an additional factor of $1 - q$. Theorem 3.1 implies that $\widetilde{L}_v$ has coefficients in $\mathbb{Z}[q, t]$, whereas the coefficients of L_v are elements of $\mathbb{Z}[[q, t]]$. We conjecture that the normalized link symmetric function $\widetilde{L}_v$ is closely connected to the Macdonald eigenoperators ∇ and Δ.

The modified Macdonald polynomials $\widetilde{H}_\mu$ form a basis for the ring of symmetric functions with coefficients in $\mathbb{Q}(q, t)$. They can be defined via triangularity relations or combinatorially [11,9]. Given a partition μ, let B_μ be the alphabet of monomials $q^i t^j$ where (i, j) ranges over the coordinates of the cells in the Ferrers diagram of μ. We compute an example in Fig. 4.

Given a symmetric function F in variables $x_1, x_2, \ldots$ and a set of monomials $A = \{a_1, a_2, \ldots, a_n\}$, we let $F[A]$ be the result of setting $x_i = a_i$ for $1 \leq i \leq n$ and $x_i = 0$ for $i > n$. Then we define two operators on symmetric functions by setting, for $\mu \vdash n$,

$$\Delta_F \widetilde{H}_\mu = F\left[B_\mu\right] \widetilde{H}_\mu \tag{45}$$

$$\nabla \widetilde{H}_\mu = \Delta_{e_n} \widetilde{H}_\mu \tag{46}$$

and expanding linearly. Note that, for $\mu \vdash n$, $e_n[B_\mu]$ is simply the product of the n monomials in B_μ; we will sometime write T_μ for the product $e_n[B_\mu]$.

Conjecture 4.1.

$$\nabla p_{1^n} = \widetilde{L}_{0^n} \tag{47}$$

$$\Delta_{e_{n-1}} p_{1^n} = \sum_{\substack{v \in \{0,1\}^n \\ |v|=1}} \widetilde{L}_v \tag{48}$$

$$\widetilde{L}_{v0} = \nabla p_1 \nabla^{-1} \widetilde{L}_v. \tag{49}$$

In fact, the first two conjectures follow from the third.

Proposition 4.1. *(49) implies both (47) and (48).*

Before proving Proposition 4.1, we should mention that Eugene Gorsky first noticed that the identity

$$\sum_{a=0}^{d} \langle \nabla p_{1^n}, e_{n-d} h_d \rangle \, a^d = (1-q)^n f_{0^n}(q,t,a) \tag{50}$$

seemed to hold and communicated this observation to the author via Jim Haglund. Gorsky's conjectured identity is a special case of Conjecture 4.1. It is also interesting to note that the operator in (49) appears in the setting of the Rational Shuffle Conjecture as $-\mathbf{Q}_{1,1}$ [4].

Proof. We prove that (49) implies (47) and (48). The fact that (49) implies (47) is clear. For the second implication, consider $v \in \{0,1\}^n$ with $|v| = 1$. Say k is the unique position such that $v_k = 1$. By (47), $\widetilde{L}_{0^{k-1}} = \nabla p_{1^{k-1}}$. By definition, $\widetilde{L}_{0^{k-1}1}$ considers γ such that $\gamma_i = 0$ if and only if $i = k$. It follows that π_k cannot be involved in any dinv pairs, and that γ_k contributes no new area. Therefore

$$\widetilde{L}_{0^{k-1}1} = p_1 \nabla p_{1^{k-1}}. \tag{51}$$

Using (47) again, we get

$$\widetilde{L}_{0^{k-1}10^{n-k}} = \nabla p_{1^{n-k}} \nabla^{-1} p_1 \nabla p_{1^{k-1}}. \tag{52}$$

We define the *Macdonald Pieri coefficients* $d_{\mu,\nu}$ by

$$p_1 \widetilde{H}_\nu = \sum_{\mu \leftarrow \nu} d_{\mu,\nu} \widetilde{H}_\mu, \tag{53}$$

where the sum is over partitions μ obtained by adding a single cell to ν. Given a standard tableau τ, let $\mu^{(i)}$ be the partition obtained by taking the cells containing $1, 2, \ldots, i$ in τ. Then each $\mu^{(i+1)}$ is obtained by adding a single cell to $\mu^{(i)}$. Let d_τ denote the product of the Macdonald Pieri coefficients

$$d_\tau = d_{\mu^{(1)},\emptyset} d_{\mu^{(2)},\mu^{(1)}} \ldots d_{\mu^{(n)},\mu^{(n-1)}}. \tag{54}$$

Now we can express the right-hand side of (52) as

$$\nabla p_{1^{n-k}}\nabla^{-1}p_1\nabla\sum_{\nu\vdash k-1}\sum_{\tau\in\mathrm{SYT}(\nu)}d_\tau\widetilde{H}_\nu \tag{55}$$

$$=\nabla p_{1^{n-k}}\nabla^{-1}p_1\sum_{\nu\vdash k-1}\sum_{\tau\in\mathrm{SYT}(\nu)}d_\tau T_\nu\widetilde{H}_\nu \tag{56}$$

$$=\nabla p_{1^{n-k}}\sum_{\lambda\vdash k}\sum_{\tau\in\mathrm{SYT}(\lambda)}d_\tau B_\lambda(\tau,k)^{-1}\widetilde{H}_\lambda \tag{57}$$

where by $B_\lambda(\tau,n)$ we mean the monomial q^it^j associated to the cell containing n in τ. Completing the computation, we get

$$\sum_{\mu\vdash n}\widetilde{H}_\mu\sum_{\tau\in\mathrm{SYT}(\mu)}d_\tau\prod_{i\neq k}B_\mu(\tau,i). \tag{58}$$

Summing over all k, we obtain $\Delta_{e_{n-1}}p_{1^n}$. $\square$

As an example of our conjecture, we can use Sage to compute

$$\langle\nabla p_{1,1},p_{1,1}\rangle=1+q+t-qt. \tag{59}$$

This expression should equal $\left\langle\widetilde{L}_{00},p_{1,1}\right\rangle$ by Conjecture 4.1. To compute this inner product using Theorem 3.1, we will use the fact that the scalar product with p_{1^n} can be computed by considering only the permutations π of $\{1,2,\dots,n\}$, as discussed in Chapter 6 of [9]. We consider the barred Fubini words 01 and $0\overline{1}$, each of which can receive labels $\pi=12$ or 21. The corresponding diagrams are

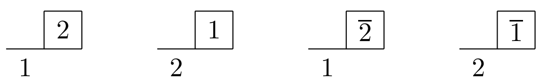

where we have moved the bars from γ_i to the corresponding π_i. The weights of these diagrams coming from Theorem 3.1 are

$$\frac{t}{1-q}\qquad\frac{1}{1-q}\qquad\frac{q}{(1-q)^2}\qquad\frac{q}{(1-q)^2} \tag{60}$$

respectively. After multiplying by the normalizing factor $(1-q)^2$ to go from L_{00} to $\widetilde{L}_{00}$, we sum the resulting weights to get

$$(1-q)t+1-q+q+q=1+q+t-qt \tag{61}$$

as desired.

After reading an earlier version of this paper, François Bergeron contacted the author with the following additional conjectures.

Conjecture 4.2 *(Bergeron, 2016).*

$$L_{v0} = L_{1v} + qL_{0v} \tag{62}$$

$$L_{0^n} = \sum_{v\in\{0,1\}^k} q^{n-|v|} L_{v0^{n-k}} \tag{63}$$

$$t\left(L_{u011v} - L_{u101v}\right) = L_{u101v} - L_{u110v} \tag{64}$$

$$\widetilde{L}_{0^a 1^b 0^c} = \nabla p_{1^c} \nabla^{-1} \widetilde{H}_{1^b} \nabla p_{1^a} \tag{65}$$

$$L_{1^a 0 1^b} = \frac{t^a - 1}{t^{a+b} - 1}\left[\nabla p_1 \nabla^{-1}, \widetilde{H}_{1^{a+b}}\right] + \widetilde{H}_{1^{a+b}} p_1 \tag{66}$$

where the bracket represents the Lie bracket and operators are applied to 1 if nothing is explicitly specified. Bergeron also observed that $L_v(x; q, 1+t)$ is e-positive. (For more context on this last statement, see Section 4 of [2].)

It is clear that (62) implies (63). We do not know of any other relations between these conjectures. We close with two more open questions.

(1) Is there a Macdonald eigenoperator expression for $\widetilde{L}_v$ for other v? Perhaps we can use ideas from the Rational Shuffle Conjecture [4], recently proved by Mellit [16].
(2) Can we generalize our conjecture for ∇p_{1^n} to "interpolate" between our conjecture and the Shuffle Theorem [5], or maybe the Square Paths Theorem [18]?

Acknowledgments

The author would like to Ben Elias and Matt Hogancamp for their exciting paper and for use of Fig. 1; Lyla Fadali for reading an earlier draft; Jim Haglund for editing and feedback; Eugene Gorsky for his comments and for the idea that Elias and Hogancamp's work could be related to Macdonald polynomials; and François Bergeron for Conjecture 4.2 along with other helpful suggestions.

References

[1] J. Alexander, A lemma on a system of knotted curves, Proc. Natl. Acad. Sci. USA 9 (1923) 93–95.
[2] F. Bergeron, Open Questions for operators related to Rectangular Catalan Combinatorics, arXiv:1603.04476, March 2016.
[3] F. Bergeron, A.M. Garsia, M. Haiman, G. Tesler, Identities and positivity conjectures for some remarkable operators in the theory of symmetric functions, Methods Appl. Anal. 6 (3) (1999) 363–420.
[4] F. Bergeron, A. Garsia, E.S. Leven, G. Xin, Compositional (km, kn)-shuffle conjectures, Int. Math. Res. Not. IMRN (October 2015).
[5] E. Carlsson, A. Mellit, A proof of the shuffle conjecture, arXiv:1508.06239 [math], August 2015.
[6] B. Elias, Matthew Hogancamp, On the computation of torus link homology, arXiv:1603.00407, March 2016.
[7] E. Gorsky, A. Negut, Refined knot invariants and Hilbert schemes, J. Math. Pures Appl. 9 (104) (2015) 403–435.

[8] E. Gorsky, A. Oblomkov, J. Rasmussen, V. Shende, Torus knots and the rational DAHA, Duke Math. J. 163 (14) (2014) 2709–2794.

[9] J. Haglund, The q, t-Catalan Numbers and the Space of Diagonal Harmonics, University Lecture Series, vol. 41, Amer. Math. Soc., 2008.

[10] J. Haglund, The combinatorics of knot invariants arising from the study of Macdonald polynomials, in: A. Beveridge, J.R. Griggs, L. Hogben, G. Musiker, P. Tetali (Eds.), Recent Trends in Combinatorics, in: The IMA Volumes in Math. and Its Applications, 2016, pp. 579–600.

[11] J. Haglund, M. Haiman, N. Loehr, A combinatorial formula for Macdonald polynomials, J. Amer. Math. Soc. 18 (2005) 735–761.

[12] M. Hogancamp, Categorified Young symmetrizers and stable homology of torus links, arXiv:1505.08148, May 2015.

[13] Mikhail Khovanov, Triply-graded link homology and Hochschild homology of Soergel bimodules, Internat. J. Math. 18 (8) (2007) 869–885.

[14] Mikhail Khovanov, Lev Rozansky, Matrix factorizations and link homology, Fund. Math. 199 (1) (2008) 1–91.

[15] A. Lascoux, B. Leclerc, J.-Y. Thibon, Ribbon tableaux, Hall–Littlewood functions, quantum affine algebras, and unipotent varieties, J. Math. Phys. 38 (2) (1997) 1041–1068.

[16] A. Mellit, Toric braids and (m, n)-parking functions, arXiv:1604.07456, April 2016.

[17] Q. Ren, Ordered partitions and drawings of rooted plane trees, Discrete Math. 338 (2015) 1–9.

[18] E. Sergel Leven, A proof of the square paths conjecture, arXiv:1601.06249, January 2016.

[19] N.J.A. Sloane, The on-line encyclopedia of integer sequences, Published electronically at http://oeis.org.

[20] R.P. Stanley, Enumerative Combinatorics, vol. 2, Cambridge University Press, 1999.