Characterizing partition functions of the edge-coloring model by rank growth

CrossMark

Alexander Schrijver [1,2]

University of Amsterdam and CWI, Amsterdam, Netherlands

A B S T R A C T

We characterize which graph invariants are partition functions of an edge-coloring model over $\mathbb{C}$, in terms of the rank growth of associated 'connection matrices'.

© 2015 Elsevier Inc. All rights reserved.

1. Introduction

Let $\mathcal{G}$ denote the collection of all undirected graphs, two of them being the same if they are isomorphic. In this paper, all graphs are finite and may have loops and multiple edges. Let $k \in \mathbb{N}$ and let $\mathbb{F}$ be a commutative ring. Call any function $y : \mathbb{N}^k \to \mathbb{F}$ a (*k-color*) *edge-coloring model* (*over* $\mathbb{F}$). In the case where y is symmetric under the action of S_k, an edge-coloring model is called a 'vertex model' by de la Harpe and Jones [5], where colors are called 'states'. The more general model was considered in the context of Holant functions by L.G. Valiant (cf. [1]).

The *partition function* of an edge-coloring model y is the function $p_y : \mathcal{G} \to \mathbb{F}$ defined for any graph $G = (V, E)$ by

$$p_y(G) := \sum_{\kappa : E \to [k]} \prod_{v \in V} y_{\kappa(\delta(v))}.$$

Here $\delta(v)$ is the set of edges incident with v. Then $\kappa(\delta(v))$ is a multisubset of $[k]$, which we identify with its incidence vector in $\mathbb{N}^k$. Moreover, we use $\mathbb{N} = \{0, 1, 2, \ldots\}$ and for $n \in \mathbb{N}$,

$$[n] := \{1, \ldots, n\}.$$

We can visualize κ as a coloring of the edges of G and $\kappa(\delta(v))$ as the multiset of colors 'seen' from v. The edge-coloring model was considered by de la Harpe and Jones [5] as a physical model, where vertices serve as particles, edges as interactions between particles, and colors as states or energy levels. It extends the Ising–Potts model. Several graph parameters are partition functions of some edge-coloring model, like the number of matchings. There exist real-valued graph parameters that are partition functions of an edge-coloring model over $\mathbb{C}$, but not over $\mathbb{R}$. (A simple example is $(-1)^{|E(G)|}$.)

In this paper, we characterize which functions $f : \mathcal{G} \to \mathbb{C}$ are the partition function of an edge-coloring model over $\mathbb{C}$. The characterization differs from an earlier characterization given in [3] (which our present characterization uses) in that it is based on the rank growth of associated 'connection matrices'.

To describe it, we need the notion of a k-fragment. For $k \in \mathbb{N}$, a *k-fragment* is an undirected graph $G = (V, E)$ together with an injective 'label' function $\lambda : [k] \to V$, where $\lambda(i)$ is a vertex of degree 1, for each $i \in [k]$. (You may alternatively view these degree-1 vertices as ends of 'half-edges', or rather of 'edge pieces', as both ends of an edge might be labeled.)

If G and H are k-fragments, the graph $G \cdot H$ is obtained from the disjoint union of G and H by identifying equally labeled vertices and by ignoring each of the k identified points as vertex, joining its two incident edges into one edge.
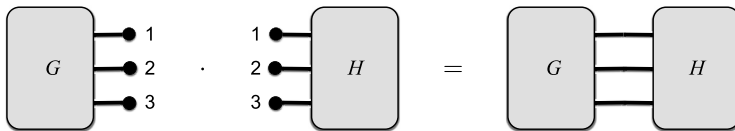
The multiplication $G \cdot H$.

(A good way to imagine this is to see a graph as a topological 1-complex.) Note that it requires (as in [8]) that we also should consider the 'vertexless loop' as possible edge of a graph, as we may create it in $G \cdot H$. We denote this vertexless loop by $\bigcirc$. Observe that if y is an edge-coloring model over $\mathbb{C}$ with n colors, then $p_y(\bigcirc) = n$.

Let $\mathcal{G}_k$ denote the collections of k-fragments. For any $f : \mathcal{G} \to \mathbb{C}$ and $k \in \mathbb{N}$, the *kth connection matrix* is the $\mathcal{G}_k \times \mathcal{G}_k$ matrix $C_{f,k}$ defined by

$$(C_{f,k})_{G,H} := f(G \cdot H)$$

for $G, H \in \mathcal{G}_k$.

Now we can formulate our characterization:

Theorem. *A function $f : \mathcal{G} \to \mathbb{C}$ is the partition function of an edge-coloring model over $\mathbb{C}$ if and only if $f(\emptyset) = 1$, $f(\bigcirc) \in \mathbb{R}$, and*

$$\mathrm{rank}(C_{f,k}) \leq f(\bigcirc)^k$$

for each $k \in \mathbb{N}$.

Let us relate this to Szegedy's theorem [8], which characterizes the partition functions of edge-coloring models over $\mathbb{R}$. Call a function $f : \mathcal{G} \to \mathbb{C}$ *multiplicative* if $f(\emptyset) = 1$ and $f(G \dot\cup H) = f(G)f(H)$ for all graphs G and H, where $G \dot\cup H$ denotes the disjoint union of G and H. Then Szegedy's theorem reads:

> A function $f : \mathcal{G} \to \mathbb{R}$ is the partition function of an edge-coloring model over $\mathbb{R}$ if and only if f is multiplicative and $C_{f,k}$ is positive semidefinite for each k.

For related results for the 'spin model' see [4] and [7].

The proof of our theorem is based on some elementary results from the representation theory of the symmetric group, and on the following alternative characterization of partition functions of edge-coloring models given in [3], which uses the Nullstellensatz and the First and Second Fundamental Theorems of Invariant theory for $O(n)$.

For any graph $G = (V, E)$, any $U \subseteq V$, and any $s : U \to V$, define

$$E_s := \{us(u) \mid u \in U\} \text{ and } G_s := (V, E \cup E_s),$$

where $us(u)$ denotes an edge connecting u and $s(u)$ (adding multiple edges if E_s intersects E). Let S_U be the group of permutations of U. Then [3]:

> A function $f : \mathcal{G} \to \mathbb{C}$ is the partition function of some k-color edge-coloring model over $\mathbb{C}$ if and only if f is multiplicative and for each graph $G = (V, E)$, each $U \subseteq V$ with $|U| = k+1$, and each $s : U \to V$: $\displaystyle\sum_{\pi \in S_U} \mathrm{sgn}(\pi) f(G_{s \circ \pi}) = 0.$ 　　(1)

2. Some results on the symmetric group

In the proof of our theorem we will need Proposition 3 below, which we prove in a number of steps. (The result might be known, and must not be a difficult exercise for those familiar with the representation theory of the symmetric group, but we did not find an explicit reference.)

We recall a few standard results from the representation theory of the symmetric group S_n (cf. James and Kerber [6]). Basis is the one-to-one relation between the partitions λ of n and the irreducible representations r_λ of S_n. Here a *partition* λ of n is a finite nonincreasing sequence $(\lambda_1, \ldots, \lambda_t)$ of positive integers with sum n. One writes $\lambda \vdash n$ if λ is a partition of n. The number t of terms of λ is called the *height* of λ, denoted by height(λ). Denote by f^λ the degree of representation r_λ (that is, the dimension of the representation space of f^λ), and by χ_λ the character of r_λ.

For any $\lambda \vdash n$, the *Young shape* Y_λ of $\lambda = (\lambda_1, \ldots, \lambda_t)$ is the following subset of $\mathbb{N}^2$:

$$Y_\lambda := \{(i, j) \mid i \in [t], j \in [\lambda_i]\}.$$

For any $\pi \in S_n$, let $o(\pi)$ denote the number of orbits of π.

Proposition 1. *For any $n \in \mathbb{N}$, $\lambda \vdash n$, and $d \in \mathbb{C}$:*

$$\sum_{\pi \in S_n} \chi_\lambda(\pi) d^{o(\pi)} = f^\lambda \prod_{(i,j) \in Y_\lambda} (d + j - i). \tag{2}$$

Proof. As both sides of (2) are polynomials in d, we can assume that $d \in \mathbb{N}$. Consider the representation r of S_n on $(\mathbb{C}^d)^{\otimes n}$ induced by

$$\pi \cdot (x_1 \otimes \cdots \otimes x_n) = x_{\pi(1)} \otimes \cdots \otimes x_{\pi(n)}$$

for $x_1, \ldots, x_n \in \mathbb{C}^d$. Note that the character χ of r satisfies $\chi(\pi) = d^{o(\pi)}$ for each $\pi \in S_n$, since, fixing a basis $e_1, \ldots, e_d$ of $\mathbb{C}^d$, it is equal to the number of $z = e_{i_1} \otimes \cdots \otimes e_{i_n}$ with $\pi \cdot z = z$; that is, it is equal to the number of $(i_1, \ldots, i_n) \in [d]^n$ with $(\pi(i_1), \ldots, \pi(i_n)) = (i_1, \ldots, i_n)$. So the requirement is that $i_j = i_k$ whenever j and k belong to the same orbit of π. Therefore, this number is equal to $d^{o(\pi)}$.

For any $\alpha \vdash n$, let μ_α be the multiplicity of r_α in r. Then, using Schur–Weyl duality,

$$\sum_{\pi \in S_n} \chi_\lambda(\pi) d^{o(\pi)} = \sum_{\pi \in S_n} \chi_\lambda(\pi) \overline{\chi(\pi)} = \sum_{\pi \in S_n} \chi_\lambda(\pi) \sum_{\alpha \vdash n} \mu_\alpha \overline{\chi_\alpha(\pi)}$$

$$= \sum_{\alpha \vdash n} \mu_\alpha \sum_{\pi \in S_n} \chi_\lambda(\pi) \overline{\chi_\alpha(\pi)} = \sum_{\alpha \vdash n} \mu_\alpha n! \delta_{\lambda, \alpha} = n! \mu_\lambda$$

$$= f^\lambda \prod_{(i,j) \in Y_\lambda} (d + j - i).$$

The last equality follows from the fact that μ_λ is equal to the degree of the irreducible representation of $\mathrm{GL}(d, \mathbb{C})$ corresponding to λ (by Schur–Weyl duality) and that this degree is equal to the last expression divided by $n!$ (cf. [2], eq. 9.28). This shows (2). $\quad\square$

For any $n \in \mathbb{N}$ and $d \in \mathbb{C}$, let $M_n(d)$ be the $S_n \times S_n$ matrix with

$$(M_n(d))_{\rho,\sigma} := d^{o(\rho\sigma^{-1})}$$

for $\rho, \sigma \in S_n$.

Proposition 2. *For any $n \in \mathbb{N}$ and $d \in \mathbb{C}$:*

$$\mathrm{rank}(M_n(d)) = \begin{cases} n! & \text{if } d \notin \mathbb{Z}, \\ \sum ((f^\lambda)^2 \mid \lambda \vdash n, \mathrm{height}(\lambda) \le |d|) & \text{if } d \in \mathbb{Z}. \end{cases} \tag{3}$$

Proof. First, we have, for any $d \in \mathbb{C}$,

$$\mathrm{rank}(M_n(-d)) = \mathrm{rank}(M_n(d)). \tag{4}$$

Indeed, note that $M_n(-d) = (-1)^n \Delta_{\mathrm{sgn}} M_n(d) \Delta_{\mathrm{sgn}}$, where Δ_{sgn} is the $S_n \times S_n$ diagonal matrix with $(\Delta_{\mathrm{sgn}})_{\pi,\pi} = \mathrm{sgn}(\pi)$ for $\pi \in S_n$. (This because $\mathrm{sgn}(\pi) = (-1)^{n-o(\pi)}$ for all π, hence $(-1)^{o(\rho\sigma^{-1})} = (-1)^n \mathrm{sgn}(\rho\sigma^{-1}) = (-1)^n \mathrm{sgn}(\rho) \mathrm{sgn}(\sigma)$.) This gives (4).

Let R be the regular representation of S_n. So, for any $\pi \in S_n$, $R(\pi)$ is the $S_n \times S_n$ matrix with

$$R(\pi)_{\rho,\sigma} = \begin{cases} 1 & \text{if } \rho = \pi\sigma, \\ 0 & \text{otherwise} \end{cases}$$

for $\rho, \sigma \in S_n$. Then

$$M_n(d) = \sum_{\pi \in S_n} d^{o(\pi)} R(\pi).$$

Now $M_n(d)$ commutes with each $R(\pi)$, and belongs to the group algebra of S_n. So $\mathrm{rank}(M_n(d))$ is equal to the sum of $(f^\lambda)^2$ taken over those λ such that $M_n(d)$ has nonzero trace in representation r_λ. That is, such that $\sum_{\pi \in S_n} \chi_\lambda(\pi) d^{o(\pi)} \ne 0$. So

$$\mathrm{rank}(M_n(d)) = \sum \left((f^\lambda)^2 \mid \lambda \vdash n, \sum_{\pi \in S_n} \chi_\lambda(\pi) d^{o(\pi)} \ne 0 \right)$$

$$= \sum ((f^\lambda)^2 \mid \lambda \vdash n, d \notin \{ i - j \mid (i,j) \in Y_\lambda \}).$$

The last equality follows from (2).

Now if $d \notin \mathbb{Z}$, then for all $\lambda \vdash n$: $d \ne i - j$ for all $(i,j) \in Y_\lambda$. So $\mathrm{rank}(M_{n,d}) = n!$. If $d \in \mathbb{Z}$, then by (4) we can assume that d is nonnegative, that is, $d \in \mathbb{N}$. Then for each $\lambda \vdash n$: $d \notin \{ i - j \mid (i,j) \in Y_\lambda \}$ if and only if $\mathrm{height}(\lambda) \le d$. This proves (3). $\quad\square$

Proposition 3. *For any $d \in \mathbb{C}$:*

$$\sup_{n \in \mathbb{N}} (\mathrm{rank}(M_n(d)))^{1/n} = \begin{cases} \infty & \text{if } d \notin \mathbb{Z}, \\ d^2 & \text{if } d \in \mathbb{Z}. \end{cases} \tag{5}$$

Proof. If $d \notin \mathbb{Z}$, the result follows directly from (3), as $\sup_n n!^{1/n} = \infty$.

If $d \in \mathbb{Z}$, then again by (4) we can assume that d is nonnegative, that is, $d \in \mathbb{N}$. Then $\mathrm{rank}(M_n(d)) \leq (d^2)^n$. Indeed, let χ be the character of the natural representation r of S_n on $(\mathbb{C}^d)^{\otimes n}$. Then $d^{o(\pi)} = \chi(\pi)$ for all $\pi \in S_n$. Hence $d^{o(\rho\sigma^{-1})} = \chi(\rho\sigma^{-1})$. So $d^{o(\rho\sigma^{-1})}$ is the trace of the product of the $d^n \times d^n$ matrices $r(\rho)$ and $r(\sigma^{-1})$. Hence $\mathrm{rank}(M_n(d)) \leq (d^n)^2 = (d^2)^n$. This proves $\leq$ in (5).

To prove the reverse inequality, consider for any $m \in \mathbb{N}$, the partition $\lambda_m = (m, \ldots, m)$ of $n := dm$, with $\mathrm{height}(\lambda_m) = d$. By the hook formula,

$$f^{\lambda_m} = n! / \prod_{i=1}^{d} \prod_{j=1}^{m} (i+j-1) = \frac{(dm)! \, 0! \, 1! \cdots (d-1)!}{m! \, (m+1)! \cdots (m+d-1)!} = \frac{(dm)!}{m!^d p(m)},$$

where (fixing d) $p(m)$ is a polynomial in m (namely $p(m) = \prod_{i=0}^{d-1} \binom{m+i}{i}$). So, by Stirling's formula, $\lim_{m \to \infty} (f^{\lambda_m})^{1/dm} = d$. By (3), we have for each m, since $\lambda_m \vdash dm$ and $\mathrm{height}(\lambda_m) = d$,

$$\mathrm{rank}(M_{dm}(d)) \geq (f^{\lambda_m})^2.$$

This gives the required inequality. $\square$

3. Proof of the theorem

Necessity being easy, we show sufficiency. As $f(\emptyset) = 1$ and $\mathrm{rank}(C_{f,0}) \leq f(\bigcirc)^0 = 1$, we know that f is multiplicative. Moreover, as $\mathrm{rank}(C_{f,1}) \leq f(\bigcirc)$, we know $f(\bigcirc) \geq 0$.

We develop some straightforward algebra. Let $k \in \mathbb{N}$. For $G, H \in \mathcal{G}_{2k}$, define the product GH as the $2k$-fragment obtained from the disjoint union of G and H by identifying vertex labeled $k+i$ in G with vertex labeled i in H, and ignoring this vertex as vertex (for $i = 1, \ldots, k$); the vertices of G labeled $1, \ldots, k$ and those of H labeled $k+1, \ldots, 2k$ make GH to a $2k$-fragment again.
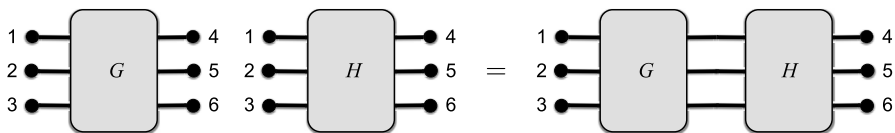
The multiplication GH.

Geometrically, one may imagine that the $2k$-fragments have the labels $1, \ldots, k$ vertically at the left and the labels $k+1, \ldots, 2k$ vertically at the right. Then GH arises by

drawing G at the left from H and connecting the right-side labels of G with the left-side labels of H, in order.

Clearly, this product is associative. Moreover, there is a unit, denoted by $\mathbf{1}_k$, consisting of k disjoint edges $e_1, \ldots, e_k$, where the ends of e_i are labeled i and $k + i$ ($i = 1, \ldots, k$).

Let $\mathbb{C}\mathcal{G}_{2k}$ be the collection of formal $\mathbb{C}$-linear combinations of elements of $\mathcal{G}_{2k}$. Extend the products $G \cdot H$ and GH bilinearly to $\mathbb{C}\mathcal{G}_{2k}$. The latter product makes $\mathbb{C}\mathcal{G}_{2k}$ to a $\mathbb{C}$-algebra.

Let $\mathcal{I}_{2k}$ be the null space of the matrix $C_{f,2k}$; that is, it consists of all $\gamma \in \mathbb{C}\mathcal{G}_{2k}$ with $f(\gamma \cdot H) = 0$ for all $H \in \mathcal{G}_{2k}$. Then $\mathcal{I}_{2k}$ is an ideal in the algebra $\mathbb{C}\mathcal{G}_{2k}$, and the quotient

$$\mathcal{A}_k := \mathbb{C}\mathcal{G}_{2k}/\mathcal{I}_{2k}$$

is an algebra of dimension $\mathrm{rank}(C_{f,2k})$. We will indicate elements of $\mathcal{A}_k$ by representatives in $\mathbb{C}\mathcal{G}_{2k}$.

Define the 'trace-like' function $\tau : \mathcal{A}_k \to \mathbb{C}$ by

$$\tau(x) := f(x \cdot \mathbf{1}_k).$$

Then $\tau(xy) = \tau(yx)$ for all $x, y \in \mathcal{A}_k$ and $\tau(\mathbf{1}_k) = f(\bigcirc)^k$. Note also that if $G, H \in \mathcal{G}_{2k}$, then $f(G \cdot H) = \tau(G\widetilde{H})$, where $\widetilde{H}$ arises from H by exchanging labels i and $k + i$ for each $i = 1, \ldots, k$. Extending this linearly to $\mathbb{C}\mathcal{G}_{2k}$, we know that for each $x \in \mathcal{A}_k$:

$$\text{if } f(x \cdot z) \neq 0 \text{ for some } z \in \mathcal{A}_k, \text{ then } \tau(xy) \neq 0 \text{ for some} \qquad (6)$$
$$y \in \mathcal{A}_k.$$

We will first show that $\mathcal{A}_k$ is semisimple. To this end, for $G \in \mathcal{G}_{2k}$ and $H \in \mathcal{G}_{2l}$, we need a product $G \sqcup H$ that can be intuitively described as: put G above H, and renumber the labels at the left hand side to $1, \ldots, k + l$ (in order) and renumber the labels at the right hand side to $k + l + 1, \ldots, 2(k + l)$ (in order).
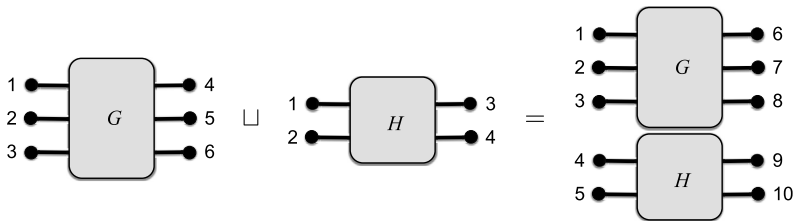
The multiplication $G \sqcup H$.

More precisely stated, $G \sqcup H$ is the $2(k+l)$-fragment obtained from the disjoint union of G and H by adding l to the labels $k + 1, \ldots, 2k$ of G, adding k to the labels $1, \ldots, l$ in H, and by adding $2k + l$ to the labels $l + 1, \ldots, 2l$ in H. This product $\sqcup$ is associative and extends bilinearly to $\mathbb{C}\mathcal{G}_{2k} \times \mathbb{C}\mathcal{G}_{2l} \to \mathbb{C}\mathcal{G}_{2(k+l)}$. Thus for $x \in \mathbb{C}\mathcal{G}_{2k}$, the m-th power $x^{\sqcup m}$ is well-defined.

Consider any $k, m \in \mathbb{N}$. For $\pi \in S_m$, let $P_{k,\pi}$ be the $2km$-fragment consisting of km disjoint edges $e_{i,j}$ for $i = 1, \ldots, m$ and $j = 1, \ldots, k$, where $e_{i,j}$ connects the vertices labeled $(i-1)k + j$ and $km + (\pi(i) - 1)k + j$.
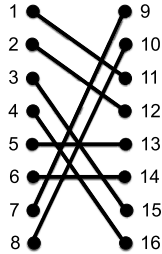
The 16-fragment $P_{2,\pi}$ with $\pi = (124) \in S_4$.

Then for any $\rho, \sigma \in S_m$ one has

$$f(x^{\sqcup m} \cdot P_{k,\pi}) = \prod_c \tau(x^{|c|}),$$

where c ranges over the orbits of π. Hence for any $\rho, \sigma \in S_m$ one has

$$f(x^{\sqcup m} P_{k,\rho} \cdot P_{k,\sigma}) = f(x^{\sqcup m} \cdot P_{k,\rho^{-1}} P_{k,\sigma}) = f(x^{\sqcup m} \cdot P_{k,\rho^{-1}\sigma}) = \prod_c \tau(x^{|c|}), \qquad (7)$$

where c ranges over the orbits of $\rho^{-1}\sigma$. We will use that, for each $x \in \mathbb{C}\mathcal{G}_{2k}$, the $S_m \times S_m$ matrix $(f(x^{\sqcup m} P_{k,\rho} \cdot P_{k,\sigma}))_{\rho,\sigma \in S_m}$ has rank at most $\mathrm{rank}(C_{f,2k})$ (since $x^{\sqcup m} P_{k,\rho}$ belongs to $\mathbb{C}\mathcal{G}_{2k}$ for each ρ).

Proposition 4. *If x is a nilpotent element of $\mathcal{A}_k$, then $\tau(x) = 0$.*

Proof. Suppose $\tau(x) \neq 0$ and x is nilpotent. Then there is a largest t with $\tau(x^t) \neq 0$. Let $y := x^t$. So $\tau(y) \neq 0$ and $\tau(y^s) = 0$ for each $s \geq 2$. By scaling we can assume that $\tau(y) = 1$.

Choose m such that $m! > f(\bigcirc)^{2km}$. By (7) we have, for any $\rho, \sigma \in S_m$,

$$f(y^{\sqcup m} P_{k,\rho} \cdot P_{k,\sigma}) = \delta_{\rho,\sigma},$$

since (7) is 0 if $\rho^{-1}\sigma$ has an orbit c with $|c| > 1$, i.e., $\rho^{-1}\sigma$ is not the identity.

So $\mathrm{rank}(C_{f,2km}) \geq m!$, contradicting the fact that $\mathrm{rank}(C_{f,2km}) \leq f(\bigcirc)^{2km} < m!$. $\square$

The following is a direct consequence of [Proposition 4](#):

Proposition 5. *$\mathcal{A}_k$ is semisimple.*

Proof. As $\mathcal{A}_k$ is finite-dimensional, it suffices to show that for each nonzero element x of $\mathcal{A}_k$ there is a y with xy not nilpotent. As $x \notin \mathcal{I}_{2k}$, we know that $f(x \cdot z) \neq 0$ for some $z \in \mathcal{A}_k$. So by (6), $\tau(xy) \neq 0$ for some $y \in \mathcal{A}_k$, and hence, by Proposition 4, xy is not nilpotent. $\square$

Proposition 6. *If x is a nonzero idempotent in $\mathcal{A}_k$, then $\tau(x)$ is a positive integer.*

Proof. Let x be any idempotent. Then for each $m \in \mathbb{N}$ and $\rho, \sigma \in S_m$, by (7):

$$f(x^{\sqcup m} P_{k,\rho} \cdot P_{k,\sigma}) = \tau(x)^{o(\rho\sigma^{-1})}.$$

So for each m:

$$\operatorname{rank}(M_m(\tau(x))) \leq \operatorname{rank}(C_{f,2km}) \leq f(\bigcirc)^{2km}.$$

Hence

$$\sup_m (\operatorname{rank}(M_m(\tau(x))))^{1/m} \leq f(\bigcirc)^{2k}.$$

By Proposition 3 this implies $\tau(x) \in \mathbb{Z}$ and $\tau(x) \leq f(\bigcirc)^k$. As $\mathbf{1}_k - x$ also is an idempotent in $\mathbb{C}\mathcal{G}_{2k}$ and as $\tau(\mathbf{1}_k) = f(\bigcirc)^k$, we have

$$f(\bigcirc)^k \geq \tau(\mathbf{1}_k - x) = f(\bigcirc)^k - \tau(x).$$

So $\tau(x) \geq 0$.

Suppose finally that x is nonzero while $\tau(x) = 0$. As $\tau(y) \geq 0$ for each idempotent y, we may assume that x is a minimal nonzero idempotent. Let J be the two-sided ideal generated by x. As $\mathcal{A}_k$ is semisimple, $J \cong \mathbb{C}^{m \times m}$ for some m. As τ is linear, there exists an $a \in J$ such that $\tau(z) = \operatorname{tr}(za)$ for each $z \in J$. As $\tau(z) = 0$ for each nilpotent z, we know that a is a diagonal matrix. As $\tau(yz) = \tau(zy)$ for all $y, z \in J$, a is in fact equal to a scalar multiple of the identity matrix.

As $x \neq 0$, $f(x \cdot z) \neq 0$ for some $z \in \mathcal{A}_k$. So by (6), $\tau(xy) \neq 0$ for some y. Hence $a \neq 0$, and so $\tau(x) \neq 0$, contradicting our assumption. $\square$

As $\mathbf{1}_1$ is an idempotent, we know that $\tau(\mathbf{1}_1)$ is a nonnegative integer, say n. So $f(\bigcirc) = n$. Let $k := n + 1$. For $\pi \in S_k$ let r_π be the $2k$-fragment consisting of k disjoint edges $e_1, \ldots, e_k$, where the ends of e_i are labeled i and $k + \pi(i)$, for $i = 1, \ldots, k$. (In fact, $r_\pi = P_{1,\pi}$ as defined above.) We define the following element q of $\mathbb{C}\mathcal{G}_{2k}$:

$$q := \sum_{\pi \in S_k} \operatorname{sgn}(\pi) r_\pi.$$

By (1) it suffices to show that $q \in \mathcal{I}_{2k}$, that is, $q = 0$ in $\mathcal{A}_k$.

Now $k!^{-1}q$ is an idempotent in $\mathbb{C}\mathcal{G}_{2k}$. Moreover,

$$\tau(q) = \sum_{\pi \in S_k} \mathrm{sgn}(\pi) n^{o(\pi)} = \sum_{\pi \in S_k} \mathrm{sgn}(\pi) \sum_{\substack{\varphi:[k] \to [n] \\ \varphi \circ \pi = \varphi}} 1 = \sum_{\varphi:[k] \to [n]} \sum_{\substack{\pi \in S_k \\ \varphi \circ \pi = \varphi}} \mathrm{sgn}(\pi) = 0,$$

since no $\varphi : [k] \to [n]$ is injective. So by Proposition 6, $q = 0$ in $\mathcal{A}_k$, as required. This finishes the proof of the theorem.

Acknowledgments

References

[1] J.-Y. Cai, P. Lu, M. Xia, Holographic algorithms by Fibonacci gates, Linear Algebra Appl. 438 (2013) 690–707.
[2] P. Cvitanović, Group Theory: Birdtracks, Lie's, and Exceptional Groups, Princeton University Press, Princeton, 2008.
[3] J. Draisma, D. Gijswijt, L. Lovász, G. Regts, A. Schrijver, Characterizing partition functions of the vertex model, J. Algebra 350 (2012) 197–206.
[4] M.H. Freedman, L. Lovász, A. Schrijver, Reflection positivity, rank connectivity, and homomorphisms of graphs, J. Amer. Math. Soc. 20 (2007) 37–51.
[5] P. de la Harpe, V.F.R. Jones, Graph invariants related to statistical mechanical models examples and problems, J. Combin. Theory Ser. B 57 (1993) 207–227.
[6] G. James, A. Kerber, The Representation Theory of the Symmetric Group, Addison–Wesley, Reading, Massachusetts, 1981.
[7] A. Schrijver, Graph invariants in the spin model, J. Combin. Theory Ser. B 99 (2009) 502–511.
[8] B. Szegedy, Edge coloring models and reflection positivity, J. Amer. Math. Soc. 20 (2007) 969–988.