The complexity of counting graph homomorphisms *

Martin Dyer and Catherine Greenhill†

Abstract

The problem of counting homomorphisms from a general graph G to a fixed graph H is a natural generalisation of graph colouring, with important applications in statistical physics. The problem of deciding whether any homomorphism exists was considered by Hell and Nešetřil. They showed that decision is NP-complete unless H has a loop or is bipartite; otherwise it is in P. We consider the problem of exactly counting such homomorphisms, and give a similarly complete characterisation. We show that counting is #P-complete unless every connected component of H is an isolated vertex without a loop, a complete graph with all loops present, or a complete unlooped bipartite graph; otherwise it is in P. We prove further that this remains true when G has bounded degree. In particular, our theorems provide the first proof of #P-completeness of the partition function of certain models from statistical physics, such as the Widom–Rowlinson model, even in graphs of maximum degree 3. Our results are proved using a mixture of spectral analysis, interpolation and combinatorial arguments.

1 Introduction

Combinatorial counting problems on graphs are important in their own right, and for their application to statistical physics. In the physics application it is often a weighted version of the problem which is of interest, corresponding to the partition function of the associated Gibbs distribution. Exactly counting proper graph colourings (i.e. evaluating the chromatic polynoial) is a classical problem and its close relative is evaluating the partition function of the Potts model in statistical physics. See, for example [24]. Here we consider the complexity of exact counting in a range of models of this type. We show that polynomial-time algorithms for exact counting are unlikely to exist, other than in certain "obvious" cases.

Many counting problems can be restated as counting the number of *homomorphisms* from the graph of interest G to a particular fixed graph H. The vertices of H correspond to colours, and the edges show which colours may be adjacent. The graph H may contain loops. Specifically, let C be a set of k colours, where k is a constant. Let $H = (C, E_H)$

*School of Computer Studies, University of Leeds, Leeds LS2 9JT, UNITED KINGDOM. Research supported by ESPRIT Working Group RAND2.

†Supported by a Leverhulme Special Research Fellowship.

be a graph with vertex set C. Given a graph $G = (V, E)$ with vertex set V, a map $X : V \mapsto C$ is called a H-colouring if

$$\{X(v), X(w)\} \in E_H \qquad \text{for all} \quad \{v, w\} \in E.$$

In other words, X is a homomorphism from G to H. Let $\Omega_H(G)$ denote the set of all H-colourings of G. The counting problem is then to determine $|\Omega_H(G)|$. We may further allow the vertices or edges of H to possess weights. The above homomorphism viewpoint has been taken previously by many authors, for example [1, 2, 13, 18, 21], and more recently by Hell and Nešetřil [11], Galluccio, Hell and Nešetřil [9] and Brightwell and Winkler [3]. The latter consider the relationship between graph homomorphisms and *phase transitions* in statistical physics.

Thus, for example, proper k-colourings of G would have H a k-clique with no loops. The Potts model corresponds to introducing a loop on every vertex and giving these loops the same positive edge-weight, the original clique edges all having (say) unit weight. The problem of counting *independent sets* corresponds to H being a single edge with one looped vertex. The (vertex) weighted version here is the well-known *hardcore lattice gas* model. The looped vertex has weight 1, and the unlooped weight $\lambda > 0$. (See, for example [15, 3, 20, 7].) Another example with physical application is the problem of counting q-particle *Widom–Rowlinson* configurations in graphs (see [25, 16]), where $q \geq 2$. This is a particular model of a gas consisting of q types of particles. The graph corresponding to q-particle Widom–Rowlinson has $q + 1$ looped vertices, one for each particle type and one representing empty sites. The latter vertex is joined to all other vertices. The graph corresponding to the 4-particle Widom–Rowlinson model is shown in Figure 1.

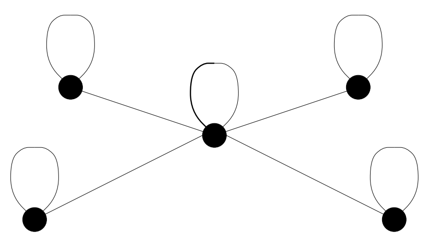

Figure 1: The graph describing 4-particle Widom–Rowlinson configurations

A further example is the Beach model [5], which is a physical system with more than one measure of maximal entropy. The graph corresponding to the Beach model is shown in Figure 2.

Hell and Nešetřil [11] gave a full characterisation of graphs H for which the corresponding decision problem (i.e. "is $\Omega_H(G)$ empty ?") is NP-complete. The decision problem corresponding to H can easily be solved in polynomial time if H has a loop or is bipartite. Conversely, Hell and Nešetřil [11] showed that if H is loopless and not bipartite then the decision problem corresponding to H is NP-complete. We will (somewhat unusually) consider the graph which consists of an unlooped isolated vertex v to

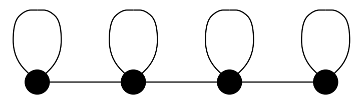

Figure 2: The graph describing the Beach model

be a complete bipartite graph with vertex bipartition $\{v\} \cup \emptyset$. This should be borne in mind when reading Hell and Nešetřil's result above and our Theorem 1.1 below.

We consider the complexity of *exactly counting* H-colourings. Denote by $\#H$ the counting problem which takes as an instance a graph G and returns the number of H-colourings of G, $|\Omega_H(G)|$. The problem $\#H$ is clearly in #P for every graph H. We will prove the following.

Theorem 1.1 *Let H be a fixed graph. The problem of counting H-colourings of graphs is #P-complete if H has a connected component which is not a complete graph with all loops present or a complete bipartite graph with no loops present. Otherwise, the counting problem is in P.*

In particular, Theorem 1.1 provides the first #P-completeness proof for counting configurations in both the Widom-Rowlinson and Beach models. Note further that #P-hardness of the partition function for any weighted version of the problem can be deduced from our results below, whenever some suitable version of the underlying counting problem is #P-complete. For example, #P-hardness of the partition functions of the appropriate weighted versions of the Widom-Rowlinson and Beach models follows easily.

In physical applications, we are usually interested in graphs of low degree. Therefore, we will prove further the following theorem.

Theorem 1.2 *Let H be a graph such that $\#H$ is #P-complete. Then there exists a constant Δ such that $\#H$ remains #P-complete when restricted to instances with maximum degree at most Δ.*

We can show that the implied Δ in Theorem 1.2 is *three* in most important cases. For example, for both the Widom-Rowlinson and Beach models this follows. Unfortunately, for a technical reason which will emerge, we are not able to assert this in general.

Theorem 1.2 shows that a counting problem which is #P-complete remains #P-complete when restricted to instances with some constant maximum degree. This is in contrast with the corresponding result for decision problems, obtained by Galluccio, Hell and Nešetřil [9]. They showed that there exist graphs H with NP-complete decision problems, such that decision is always a polynomial-time operation when restricted to graphs with maximum degree 3. Moreover, there does not seem to be any clear characterisation of graphs for which the corresponding decision problem is polynomial-time, when restricted to graphs with a given constant maximum degree. Indeed, it is possible that no such characterisation exists.

Theorem 1.1 is proved by using a mixture of algebraic and combinatorial methods. The algebraic tools involve spectral analysis and interpolation, using operations on graphs called "stretchings" and "thickenings". This has a flavour of Jaeger, Vertigan and Welsh's approach to proving #P-hardness of the Tutte polynomial [14]. In particular, we present a useful result (Lemma 3.4) which we believe is the first interpolation proof involving eigenvalues. Unfortunately, the algebraic tools do not seem sufficient, by themselves, to prove Theorem 1.1. Therefore the proof is completed by a combinatorial case analysis, somewhat analogous to Hell and Nešetřil's approach to the decision problem. However, for the decision problem it is easy to see that attention can be restricted to the case where H is connected. In the counting problem this fact is far from obvious, and forms the content of our central Theorem 4.1.

Two #P-complete problems are the starting points of all our reductions. The first is the problem of counting proper 3-colourings of graphs. (There is an easy and well-known reduction from this problem to that of counting proper q-colourings of graphs, for any $q \geq 4$.) This 3-colouring problem is #P-complete even when restricted to bipartite graphs, as shown by Linial [19]. The second "base" problem is that of counting independent sets in graphs. This problem is #P-complete even when restricted to line graphs, as shown by Valiant [23].

The plan of the paper is as follows. In the next section we introduce the operations of stretchings and thickenings. In Section 3 we use these operations to prove some interpolation results. In Section 3.1 we show how the counting problem for a singular weight matrix can always be solved by reduction from the counting problem corresponding to a certain nonsingular matrix, possibly by introducing vertex weights. In Section 3.2 we show that these vertex weights can be ignored in any proof of #P-hardness. In Section 4 we prove Theorem 1.1, by analysing cases. Finally, in Section 5 we prove Theorem 1.2, and show that all graph homomorphism counting problems are easy on instances with maximum degree 2.

1.1 Edge weights and vertex weights

The graph H is described by the $k \times k$ adjacency matrix A, where $A_{ij} = 1$ if $\{i, j\} \in E_H$ and $A_{ij} = 0$ otherwise, for all i, $j \in C$. Other problems involve weighted versions of this set-up. The weights may be on the vertices of H or on the edges. The edge weights may be stored in a symmetric matrix A, called a *weight matrix*, such that $A_{ij} = 0$ if and only of $\{i, j\} \notin E_H$. Our focus throughout the paper is on counting graph homomorphisms (where all edge weights and all vertex weights equal 1). In the proofs, however, it is usually more convenient to work with the adjacency matrix of the graph (or with more general weight matrices). Of course, the weighted problems are of interest in their own right, as they are used in statistical physics.

Suppose that the vertex weights are $\{\lambda_i\}_{i\in C}$. We assume that $\lambda_i > 0$ for all $i \in C$. Let D be the diagonal matrix with λ_i in the (i, i) position. Thus D is an invertible diagonal matrix. Let the edge weights be stored in the matrix A. If $X \in \Omega_H(G)$, let

$$w_A(X) = \prod_{\{v,w\}\in E} A_{X(v)X(w)}, \quad \widetilde{w}_D(X) = \prod_{v\in V} \lambda_{X(v)}.$$

Thus $w_A(X)$ measures the edge-weighting of the H-colouring X, and $\widetilde{w}_D(X)$ measures the vertex-weighting of X. In total, the H-colouring X has weight $w_{A,D}(X)$, defined by

$$w_{A,D}(X) = w_A(X)\,\widetilde{w}_D(X) = \prod_{\{v,w\}\in E} A_{X(v)X(w)} \prod_{v\in V} \lambda_{X(v)}.$$

There is a slight abuse in writing $A_{X(v)X(w)}$, where $\{v, w\} \in E$, as the indices of a matrix are ordered while the endpoints of an edge are not. However, as A is symmetric this does not cause any harm. We are interested in $Z_{A,D}(G)$, defined by

$$Z_{A,D}(G) = \sum_{X\in\Omega_H(G)} w_{A,D}(X).$$

If all vertex weights are equal to 1, we write $w_A(X)$ and $Z_A(G)$ instead of $w_{A,D}(X)$ and $Z_{A,D}(G)$ respectively. If in addition all edge weights are equal to 1, then $Z_A(G) = |\Omega_H(G)|$.

2 Stretchings and thickenings

Recently, most #P-completeness proofs have used interpolation as the main tool in building polynomial-time reduction. See, for example, [22, 23, 14, 4, 10]. (For some #P-completeness proofs which do not use interpolation, see [12]). These interpolations are often designed to preserve desirable properties, such as bounded maximum degree of a graph. Two tools often used in these proofs are *stretchings* and *thickenings*, which are now described in the graph setting.

Let P_r denote the *path* with $r + 1$ vertices $u_0, \dots, u_r$ and r edges $\{u_i, u_{i+1}\}$, for $0 \le i < r$, where $r \ge 1$. Let $G = (V, E)$ be a given graph. The *r-stretch* of G, denoted by S_rG, is obtained by replacing each edge $\{v, w\}$ in E by a copy of the path P_r, using the identifications $v = u_0$, $w = u_r$. We can also define the r-stretch of G with respect to a subset $F \subseteq E$ of edges of G, denoted by ${S_r}^{(F)}(G)$. To form ${S_r}^{(F)}(G)$, replace each edge in F by a copy of P_r.

We seek an expression for $Z_A(S_rG)$. For $i, j \in C$ let

$${\Omega_H}^{(i,j)}(P_r) = \{X \in \Omega_H(P_r) \mid X(u_0) = i,\;\; X(u_r) = j\}\,. \tag{1}$$

It is not difficult to see that these sets form a partition of $\Omega_H(P_r)$. Let D be a diagonal matrix of positive vertex weights $\{\lambda_i \mid i \in C\}$ and let Π be the diagonal matrix with (i, i) entry equal to $\sqrt{\lambda_i}$. Using induction, one can prove that

$$\sum_{X\in{\Omega_H}^{(i,j)}(P_r)} w_{A,D}(X) = \sqrt{\lambda_i\lambda_j}\,{B^r}_{ij} \tag{2}$$

where $B = \Pi A \Pi$. (Note that the proof follows exactly the same steps as the proof of Lemma 3.8, given below.) In the vertex-unweighted case, we have

$$\sum_{X\in{\Omega_H}^{(i,j)}(P_r)} w_A(X) = {A^r}_{ij}. \tag{3}$$

There is a relationship between the graph A^r and random walks on graphs which are endowed with H-colourings. Consider performing a random walk on the vertices of a graph G. If ${A^r}_{ij} \neq 0$ then it is possible to walk from a vertex coloured i to a vertex coloured j in exactly r steps.

Using (3), we can express $Z_A(S_r G)$ in terms of the entries of A^r, for $r \geq 1$.

Corollary 2.1 *Let $r \geq 1$. Then*

$$Z_A(S_r G) = \sum_{X:V \to C} \prod_{\{v,w\} \in E} {A^r}_{X(v)X(w)} = Z_{A^r}(G).$$

Proof. Suppose that $X \in \Omega_H(S_r G)$. We can think of X as being formed from the following ingredients: we have a map $Y : V \mapsto C$ which is the restriction of X to the vertices of G, and for each edge $\{v, w\} \in E$ we have the restriction of X to the path P_r between v and w, with endpoints coloured $Y(v)$ and $Y(w)$. This construction can be reversed, giving a bijection between $\Omega_H(S_r G)$ and

$$\bigcup_{Y:V \mapsto C} \prod_{\{v,w\} \in E} {\Omega_H}^{(Y(v),Y(w))}(P_r).$$

Using this bijection, we find that

$$\begin{aligned} Z_A(S_r G) &= \sum_{X \in \Omega_H(S_r G)} w_A(X) \\ &= \sum_{Y:V \mapsto C} \prod_{\{v,w\} \in E} \sum_{Z \in {\Omega_H}^{(Y(v),Y(w))}(P_r)} w_A(Z) \\ &= \sum_{Y:V \mapsto C} \prod_{\{v,w\} \in E} {A^r}_{Y(v)Y(w)}, \end{aligned}$$

as claimed. The second equality follows from the bijection and the third equality follows from (3). □

Think of S_r as an operation which acts on graphs. Let $\sigma_r A = A^r$, the rth power of A. Then Corollary 2.1 can be restated as

$$Z_A(S_r G) = Z_{\sigma_r A}(G).$$

We can think of σ_r as the r-stretch operation for weight matrices. Notice that S_1 and σ_1 are both identity maps. Moreover $S_r S_t = S_{rt}$ and $\sigma_r \sigma_t = \sigma_{rt}$ for all $r, t \geq 1$. The vertex-weighted form of Corollary 2.1 states that

$$Z_{A,D}(S_r G) = Z_{A\sigma_{r-1}(DA),D}(G)$$

for any graph G and any diagonal matrix D of vertex weights. This statement is proved using (2).

Now let $p \geq 1$. The *p-thickening* of the graph G, denoted by T_pG, is obtained by replacing each edge by p copies of itself. This results in a multigraph where each edge has multiplicity p. We can also define the p-thickening of a graph G with respect to a subset $F \subseteq E$ of edges of G, denoted by ${T_p}^{(F)}(G)$. Form ${T_p}^{(F)}(G)$ by replacing each edge in F by p copies of itself.

Since the endpoints of an edge of G become endpoints of p edges in T_pG, we see immediately that

$$Z_A(T_pG) = \sum_{X:V \to C} \prod_{\{v,w\} \in E} {A_{X(v)X(w)}}^p. \tag{4}$$

Think of T_p as an operation which acts on (multi)graphs. Let $\tau_p A$ denote the matrix whose (i,j) entry is equal to ${A_{ij}}^p$. Then $Z_A(T_pG) = Z_{\tau_p A}(G)$. We can think of τ_p as the p-thickening operation for weight matrices. Notice that T_1 and τ_1 are both identity maps. Moreover, $T_pT_q = T_{pq}$ and $\tau_p\tau_q = \tau_{pq}$ for all $p, q \geq 1$.

Thickenings do not interfere with vertex-weighted problems. Specifically, it is not difficult to see that

$$Z_{A,D}(T_pG) = Z_{\tau_p A, D}(G)$$

for all weight matrices A, invertible diagonal matrices D and graphs G. This gives rise to a polynomial-time reduction from $\mathrm{EVAL}(\tau_p A, D)$ to $\mathrm{EVAL}(A, D)$ for any $p \geq 1$ which is polynomially bounded.

In some circumstances, we wish to perform the stretching or thickening operation with respect to some subset of edges only. For the rest of this section, however, we consider the case where all edges are involved. Now consider how the thickening and stretching operations interact with each other.

The r-stretch of the p-thickening of G is denoted by S_rT_pG. By inspection, we see that

$$Z_A(S_rT_pG) = \sum_{X:V \to C} \prod_{\{v,w\} \in E} {A^r{}_{X(v)X(w)}}^p = Z_{\tau_p\sigma_r A}(G). \tag{5}$$

Notice that the thickening and stretching operations are applied to A in the reverse order that they are applied to G. The p-thickening of the r-stretch of G is denoted by T_pS_rG. Let $B = \tau_p A$, so that $B_{ij} = {A_{ij}}^p$. Then, by inspection, we see that

$$Z_A(T_pS_rG) = \sum_{X:V \to C} \prod_{\{v,w\} \in E} B^r{}_{X(v)X(w)} = Z_{\sigma_r\tau_p A}(G).$$

Again, the order of the stretching and thickening operations is reversed for the weight matrix. For illustration, the graphs S_5T_3e and T_3S_5e are shown in Figure 3, where e is a single edge.

We would like to be able to apply arbitrary compositions of the thickening and stretching operations.

Lemma 2.1 *Any composition of thickening and stretching operations is equivalent to a composition of the form*

$$S_{r_\ell}T_{p_\ell}S_{r_{\ell-1}}T_{p_{\ell-1}} \cdots S_{r_2}T_{p_2}S_{r_1}T_{p_1},$$

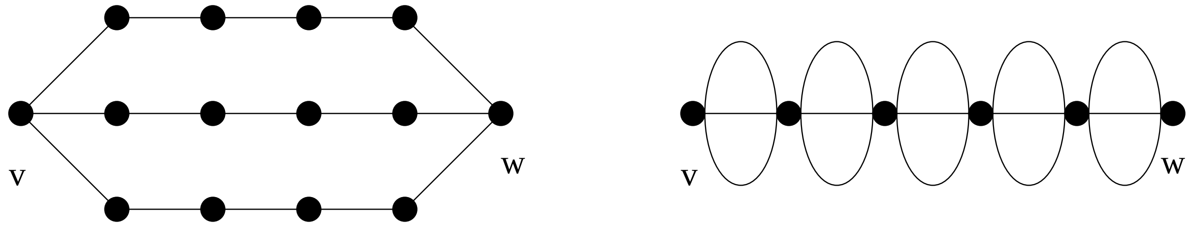

Figure 3: The graphs S_5T_3e and T_3S_5e, where $e = \{v, w\}$

where $\ell \geq 1$ and $1 \leq r_i, p_i$ for $1 \leq i \leq \ell$. In addition, we have

$$Z_A(S_{r_\ell}T_{p_\ell}S_{r_{\ell-1}}T_{p_{\ell-1}}\cdots S_{r_1}T_{p_1}G) = Z_{\tau_{p_1}\sigma_{r_1}\cdots\tau_{p_{\ell-1}}\sigma_{r_{\ell-1}}\tau_{p_\ell}\sigma_{r_\ell}A}(G).$$

Proof. This first statement follows since $S_rS_t = S_{rt}$, $T_pT_q = T_{pq}$, $S_1 = \mathrm{id}$ and $T_1 = \mathrm{id}$ for all r, s, p, $q \geq 1$. The second statement can easily be proved by induction on ℓ. □

3 Interpolation results

We now describe some interpolation results which are obtained using stretchings and thickenings. First, let us introduce some notation. Let H be a graph and A a weight matrix on H. Let D be a nonsingular diagonal matrix of vertex weights. The problems $\#H$, EVAL(A) and EVAL(A, D) are defined below.

PROBLEM:	$\#H$	EVAL(A)	EVAL(A, D)
INSTANCE:	A graph G	A graph G	A graph G
OUTPUT:	$\lvert\Omega_H(G)\rvert$	$Z_A(G)$	$Z_{A,D}(G)$

If A is the adjacency matrix of the graph H, then the problems $\#H$ and EVAL(A) are identical. Similarly, if D is the identity matrix then the problems EVAL(A) and EVAL(A, D) are identical.

The operations of stretching and thickening give rise to polynomial-time reductions, as below.

Lemma 3.1 *Suppose that r, p are positive integer constants and A is a weight matrix. There is a polynomial-time reduction from* EVAL($\tau_p\sigma_rA$) *to the problem* EVAL(A).

Proof. Let G be an instance of EVAL(A). We can form the graph S_rT_pG from G in polynomial time, and $Z_{\tau_p \sigma_r A}(G) = Z_A(S_r T_p G)$ by (5). This completes the polynomial-time reduction. $\square$

This result can be extended to constant length compositions of stretching and thickening operations, using Lemma 2.1. Something slightly more complicated can be said if vertex weights are in use.

Many polynomial-time reductions involve the following standard interpolation technique, as used in [22, 4]. Although the result is well-known, for completeness we include a proof.

Lemma 3.2 *Let $w_1, \ldots, w_r$ be known distinct nonzero constants. Suppose that we know values $f_1, \ldots, f_r$ such that*

$$f_s = \sum_{i=1}^{r} c_i \, {w_i}^s$$

for $1 \leq s \leq r$. The coefficients $c_1, \ldots, c_r$ can be evaluated in polynomial time.

Proof. We can express the equations in matrix form, as

$$\begin{bmatrix} f_1 \\ f_2 \\ \vdots \\ f_r \end{bmatrix} = \begin{bmatrix} w_1 & w_2 & \cdots & w_r \\ {w_1}^2 & {w_2}^2 & \cdots & {w_r}^2 \\ \vdots & \vdots & \ddots & \vdots \\ {w_1}^r & {w_2}^r & \cdots & {w_r}^r \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \\ \vdots \\ c_r \end{bmatrix}.$$

The $r \times r$ matrix in the above equation is invertible. To see this, divide each column by the entry in the 1st row of that column. The result is a Vandermonde matrix with distinct columns. This matrix is invertible, hence so is the original matrix. We can invert the matrix in polynomial time, to solve for $c_1, \ldots, c_r$. $\square$

The next two results concern the set of distinct weights of H-colourings of a given graph. In order to describe this result, we need some notation. Let $\mathcal{U}_A(G)$ be the set of all distinct weights of H-colourings of G; that is,

$$\mathcal{U}_A(G) = \{ w_A(X) \mid X \in \Omega_H(G) \} .$$

Suppose that G has m edges and H has k vertices. Let the nonzero entries of A be $\mu_1, \ldots, \mu_s$. Then $\mathcal{U}_A(G) \subseteq \mathcal{W}_m(A)$, where

$$\mathcal{W}_m(A) = \{ {\mu_1}^{c_1} \cdots {\mu_s}^{c_s} \mid 0 \leq c_i \leq m \ \text{ for } \ 1 \leq i \leq s, \ \ \mu_1 + \cdots + \mu_s = m \} .$$

It is not difficult to show that

$$|\mathcal{U}_A(G)| \leq |\mathcal{W}_m(A)| \leq (m+1)^{k^2}. \tag{6}$$

Let $N_A(G, w)$ be the number of H-colourings of G with weight equal to w, for any real number w. That is,

$$N_A(G, w) = | \{ X \in \Omega_H(G) \mid w_A(X) = w \} |.$$

Clearly $N_A(G, w) = 0$ if $w \notin \mathcal{U}_A(G)$. For any real number w, denote by $\mathrm{EVAL}(A, w)$ the counting problem which takes as instance a graph G and returns the output $N_A(G, w)$.

Lemma 3.3 *Let A be a weight matrix for the graph H. There is a polynomial-time reduction from the problem* EVAL(A, w) *to the problem* EVAL(A)*, for every real number w.*

Proof. The set $\mathcal{W}_m(A)$, defined above, can be constructed explicitly in polynomial time and contains only nonzero entries. Suppose that $\mathcal{W}_m(A) = \{w_1, \ldots, w_t\}$ so that there are t distinct elements in $\mathcal{W}_m(A)$. If $w \notin \mathcal{W}_m(A)$ then $N_A(G, w) = 0$. Assume now that $w = w_j \in \mathcal{W}_m(A)$. For $1 \le p \le t$ form the graph T_pG, the p-thickening of G. By (6), the integer t is polynomially bounded. For each value of p, the graph T_pG can be formed in polynomial time. Therefore we can form all t graphs in polynomial time. Using (4), we see that

$$Z_A(T_pG) = \sum_{X:V\to C} \prod_{\{v,w\}\in E} A_{X(v)X(w)}{}^p = \sum_{w\in\mathcal{W}_m(A)} w^p\, N_A(G, w).$$

The values in $\mathcal{W}_m(A)$ are known, distinct and nonzero. Using Lemma 3.2, we can calculate the coefficients $N_A(G, w)$ for all $w \in \mathcal{W}_m(A)$, in polynomial time. In particular, we know $N_A(G, w_j)$, the quantity of interest. This completes the polynomial-time reduction. □

Corollary 3.1 *Let H be a graph and let A be a weight matrix for H. There is a polynomial-time reduction from $\#H$ to* EVAL(A)*.*

Proof. Using the reduction of Lemma 3.3, we obtain the value of $N_A(G, w)$ for all $w \in \mathcal{W}_m(A)$. Summing these values, we obtain $|\Omega_H(G)|$. □

Corollary 3.2 *Suppose that the distinct entries of A are pairwise coprime positive integers, (where 1 is considered to be coprime to all integers). Let $\mathcal{S}$ be some subset of these entries. Let B be the matrix obtained by replacing all entries of A by 1 if they belong to $\mathcal{S}$, and replacing them by 0 otherwise. Then there is a polynomial-time reduction from the problem* EVAL(B) *to the problem* EVAL(A)*.*

Proof. Write $\{\mu_1, \ldots, \mu_r\}$ for the set of distinct entries of A. Let G be a given graph with m edges. The set $\mathcal{W}_m(A)$ can be written as

$$\mathcal{W}_m(A) = \left\{ {\mu_1}^{\alpha_1} \cdots {\mu_r}^{\alpha_r} \mid \ 0 \le \alpha_i \le m \ \text{ for } \ 1 \le i \le r, \ \sum_{i=1}^{r} \alpha_i = m \right\}.$$

Since the μ_i are coprime, these weights are distinct. That is, if the equation

$${\mu_1}^{\alpha_1} \cdots {\mu_r}^{\alpha_r} = {\mu_1}^{\beta_1} \cdots {\mu_r}^{\beta_r}$$

holds for two elements of $\mathcal{W}_m(A)$, then $(\alpha_1, \ldots, \alpha_r) = (\beta_1, \ldots, \beta_r)$. Define the subset $\mathcal{V}_m(\mathcal{S}, A)$ of $\mathcal{W}_m(A)$ by

$$\mathcal{V}_m(\mathcal{S}, A) = \{ {\mu_1}^{\alpha_1} \cdots {\mu_r}^{\alpha_r} \in \mathcal{W}_m(A) \mid \alpha_i = 0 \ \text{ whenever } \ \mu_i \notin \mathcal{S} \}.$$

That is, $\mathcal{V}_m(\mathcal{S}, A)$ is the set of all candidate weights which are coprime to all entries of $A \setminus \mathcal{S}$. It is not difficult to see that

$$Z_B(G) = \sum_{w \in \mathcal{V}_m(\mathcal{S},A)} N_A(G, w).$$

The set $\mathcal{V}_m(\mathcal{S}, A)$ can be computed in polynomial time. This completes the polynomial-time reduction. □

Corollary 3.3 *Let A be a symmetric matrix with every entry a nonnegative integer, and let B be the matrix obtained by replacing all the entries of A which do not equal 1 by zero. Then there is a polynomial-time reduction from* EVAL(B) *to* EVAL(A).

Proof. Using the reduction of Lemma 3.3, we can obtain the value of $N_A(G, 1)$. But $w_A(X) \neq 1$ unless all edges in G are given weight 1 by the H-colouring X. This is only possible if none of the edge-weights with greater than 1 are used. Hence $N_A(G, 1) = Z_B(G)$, as required. □

Let $G = (V, E)$ be a graph and let A a weight matrix. For $F \subseteq E$, define

$$w_A{}^{(F)}(X) = \prod_{\{v,w\} \in F} A_{X(v)X(w)}.$$

Thus $w_A(X) = w_A{}^{(E)}(G)$. The following lemma is very useful, and is proved using interpolation involving the *eigenvalues* of the matrix A. To the best of our knowledge, this is the first result involving interpolation on eigenvalues.

Lemma 3.4 *Let A be a nonsingular symmetric matrix, and let G be a given graph. Let $F \subseteq E$ be a subset of edges of G and let $m = |F|$. Suppose that we know the values of*

$$f_r(G) = \sum_{X:V \mapsto C} c_A(X)\, w_{\sigma_r A}{}^{(F)}(X)$$

for $1 \leq r \leq (m+1)^k$, where c_A is any function which depends on A but not on r. Then we can evaluate

$$\sum_{X:V \mapsto C} c_A(X)\, w_I{}^{(F)}(X)$$

in polynomial time, where I is the $k \times k$ identity matrix.

Proof. Let the eigenvalues of A be $\alpha_1, \ldots, \alpha_k$. These eigenvalues can be found computationally in polynomial time (indeed, in constant time, since k is a constant). Let L be the diagonal matrix such that $L_{ii} = \alpha_i$ for $1 \le i \le k$. Since the matrix A is symmetric, there exists an orthogonal matrix Q such that $QLQ^T = A$. Now

$$\begin{aligned}
f_r(G) &= \sum_{X:V\mapsto C} c_A(X) \prod_{\{v,w\}\in F} {A^r}_{X(v)X(w)} \\
&= \sum_{X:V\mapsto C} c_A(X) \prod_{\{v,w\}\in F} (Q\, L^r\, Q^T)_{X(v)X(w)} \\
&= \sum_{X:V\mapsto C} c_A(X) \prod_{\{v,w\}\in F} \sum_{\ell\in C} Q_{X(v)\ell}\, Q_{X(w)\ell}\, (\alpha_\ell)^r. \qquad (7)
\end{aligned}$$

Let $\mathcal{S}$ be defined by

$$\mathcal{S} = \{ {\alpha_1}^{c_1} \cdots {\alpha_k}^{c_k} \mid 0 \le c_i \le m \text{ for } 1 \le i \le k,\; c_1 + \cdots + c_k = m \}.$$

Since we know the eigenvalues of A explicitly, the set $\mathcal{S}$ can be constructed explicitly in polynomial time. We can write

$$f_r(G) = \sum_{s\in S} a_s\, s^r \qquad (8)$$

for $1 \le r \le (m+1)^k$, where the a_s are some (unknown) coefficients. The elements of $\mathcal{S}$ are known, distinct and nonzero. Therefore, we can obtain the coefficients a_s for $s \in \mathcal{S}$ in polynomial time, using Lemma 3.2. Thus we can calculate $Y = \sum_{s\in\mathcal{S}} a_s$. But Y is obtained by setting $r = 0$ in (8). Hence Y is also equal to the value obtained by setting $r = 0$ in (7), since both (7) and (8) are equations for $f_r(G)$. Therefore

$$\begin{aligned}
Y &= \sum_{X:V\mapsto C} c_A(X) \prod_{\{v,w\}\in F} \sum_{\ell\in C} Q_{X(v)\ell}\, Q_{X(w)\ell} \\
&= \sum_{X:V\mapsto C} c_A(X) \prod_{\{v,w\}\in F} (Q\, Q^T)_{X(v)X(w)} \\
&= \sum_{X:V\mapsto C} c_A(X)\, {w_I}^{(F)}(X).
\end{aligned}$$

This completes the proof. □

3.1 Singular weight matrices

In order to apply Lemma 3.4 we need a nonsingular symmetric matrix A. We now give a series of lemmas which show how to proceed when the adjacency matrix A is singular.

Lemma 3.5 *Let A be a symmetric 0-1 matrix which has a pair of linearly dependent columns. Then there exists a symmetric 0-1 matrix A' with no two linearly dependent columns, and a positive diagonal matrix D, such that the problems* EVAL(A', D) *and* EVAL(A) *are equivalent. Moreover, the matrices A' and D can be constructed from A in constant time.*

Proof. Suppose that A is the adjacency matrix of the graph H. Let H' be the subgraph of H which is formed from H as follows. First, delete all isolated vertices. Next, suppose that $\{v_1, \ldots, v_\ell\}$ are vertices of H which have the same set of neighbours as each other. Delete $v_2, \ldots, v_\ell$ from H. Continue until no two vertices have the same set of neighbours, and call the resulting graph H'. Let A' be the adjacency matrix of H'. Then A' has no zero columns or repeated columns. Finally, let D be the diagonal matrix with the same number of rows and columns as A', such that D_{ii} equals the number of vertices in H which have the same neighbours as the ith vertex of H' (considered as a vertex in H). It is not difficult to see that $\mathrm{EVAL}(A)$ and $\mathrm{EVAL}(A', D)$ are equivalent, and that A' and D can be formed from A in constant time. □

The following will be helpful when performing 2-stretches.

Lemma 3.6 *Let A be a symmetric matrix, and D be an invertible diagonal matrix. If ADA has two linearly dependent columns then A has two linearly dependent columns.*

Proof. Let $Q = \Pi A$ where $\Pi = D^{1/2}$. Then $ADA = Q^T Q$. Let a_i, q_i denote the ith column of A, Q respectively. Suppose that A has no two linearly dependent columns. It follows that Q has no two linearly dependent columns. Therefore, by the Cauchy–Schwarz inequality,

$$q_i{}^T q_j < \sqrt{(q_i{}^T q_i)(q_j{}^T q_j)}$$

whenever $i \neq j$. Fix i, j such that $1 \leq i < j \leq k$. The ith and jth columns of ADA contain a nonsingular 2×2 submatrix

$$\begin{bmatrix} q_i{}^T q_i & q_i{}^T q_j \\ q_i{}^T q_j & q_j{}^T q_j \end{bmatrix}.$$

Therefore these columns are not linearly dependent. □

We now establish an upper bound on the off-diagonal of ADA, when A has no two linearly dependent columns.

Lemma 3.7 *Suppose that A is a symmetric 0-1 matrix with no two linearly dependent columns. Let D be a diagonal matrix such that every diagonal entry $D_{ii} = \lambda_i$ is positive. Define $\lambda_{\min} = \min\{\lambda_i \mid i \in C\}$, and let*

$$\gamma = \exp\left(-\frac{\lambda_{\min}}{2\mathrm{Tr}(D)}\right),$$

where $\mathrm{Tr}(D) = \sum_{i \in C} \lambda_i$. Then

$$(ADA)_{ij} \leq \gamma \sqrt{(ADA)_{ii}(ADA)_{jj}}$$

for all $i \neq j$.

Proof. As in the proof of Lemma 3.6, let $\Pi = \sqrt{D}$ and let $Q = \Pi A$. Then $ADA = Q^T Q$. Let q_i denote the ith column of q. Define the sets N_i by $N_i = \{\ell \in C \mid A_{\ell i} \neq 0\}$ for all $i \in C$. Then

$$q_i{}^T q_j = \sum_{\ell \in N_i \cap N_j} \lambda_\ell$$

for all $i, j \in C$. Now let i and j be fixed, distinct elements of C. Since A is a 0-1 matrix with no repeated rows, the sets N_i and N_j must be different. Hence, without loss of generality, $N_i \cap N_j$ is a strict subset of N_i. Therefore

$$q_i{}^T q_j \leq \sum_{\ell \in N_i} \lambda_\ell - \lambda_{\min} \leq q_i{}^T q_i - \lambda_{\min}.$$

It follows that

$$(q_i{}^T q_j)^2 \leq \left(q_i{}^T q_i - \lambda_{\min}\right) q_i{}^T q_j$$

and so

$$\begin{aligned} \frac{(q_i{}^T q_j)^2}{(q_i{}^T q_i)(q_j{}^T q_j)} &\leq 1 - \frac{\lambda_{\min}}{q_i{}^T q_i} \\ &\leq 1 - \frac{\lambda_{\min}}{\mathrm{Tr}(D)} \\ &\leq e^{-\lambda_{\min}/\mathrm{Tr(D)}}. \end{aligned}$$

as required. $\square$

We can now prove the main result of this section, showing that the p-thickening of ADA is nonsingular when p is large enough, whenever A has no two linearly dependent columns.

Theorem 3.1 *Let A be a singular weight matrix with no two linearly dependent columns. Let D be a diagonal matrix of positive vertex weights. Let $\lambda_{\min} = \min\{\lambda_i \mid i \in C\}$. The matrix $B = \tau_p(ADA)$ is nonsingular, where*

$$p \geq \lceil 2\mathrm{Tr}(D) \log(2k)/\lambda_{\min} \rceil + 1. \tag{9}$$

Proof. Let $A' = ADA$. By Lemma 3.6, ADA has no two linearly dependent columns. We show that the p-thickening of A' is nonsingular for sufficiently large values of p. Specifically, we prove that the value of p quoted above is large enough.

Consider the determinant of A'. Each term of $\det(A')$ has the form

$$\pm \prod_{i=1}^{k} A'_{i\theta(i)},$$

where θ is a permutation on $\{1,\ldots,k\}$. Let γ be as defined in Lemma 3.7 and let $t(\theta) = |\,\{i \mid \theta(i) \neq i\}\,|$. Then, using Lemma 3.7,

$$\prod_{i=1}^{k} {A'}_{i\theta(i)} \leq \gamma^{t(\theta)} \prod_{i=1}^{k} \sqrt{{A'}_{ii}} \prod_{i=1}^{k} \sqrt{{A'}_{\theta(i)\theta(i)}} = \gamma^{t(\theta)} \prod_{i=1}^{k} {A'}_{ii},$$

with equality holding if and only if $t(\theta) = 0$.

Suppose that $p \geq 1$, and consider the p-thickening of A'. Each term of $\det(\tau_p A')$ has the form

$$\pm \prod_{i=1}^{k} {{A'}_{i\theta(i)}}^{p}$$

for some permutation θ. Now

$$|\,\{\theta \in \mathrm{Sym}(k) \mid t(\theta) = t\}\,| \leq \binom{k}{t} t! \leq k^t$$

for $0 \leq t \leq k$. By separating out the identity permutation, and subtracting all other terms, we find that

$$\begin{aligned}
\det(\tau_p A') &\geq \left(\prod_{i=1}^{k} {A'}_{ii}\right)^p - \left(\prod_{i=1}^{k} {A'}_{ii}\right)^p \sum_{t=1}^{k} k^t\, \gamma^{pt} \\
&> \left(\prod_{i=1}^{k} {A'}_{ii}\right)^p \left(1 - \frac{k\gamma^p}{1 - k\gamma^p}\right) \\
&= \left(\prod_{i=1}^{k} {A'}_{ii}\right)^p \frac{1 - 2k\gamma^p}{1 - k\gamma^p}.
\end{aligned}$$

This quantity is positive whenever $2k\gamma^p < 1$. Rearranging, we find that this holds whenever $p > \lceil 2\mathrm{Tr}(D) \log(2k)/\lambda_{\min} \rceil$. Therefore $\tau_p A'$ is a nonsingular matrix for $p \geq \lceil 2\mathrm{Tr}(D) \log(2k)/\lambda_{\min} \rceil + 1$. This completes the proof. □

Note that, in the vertex-unweighted case, it suffices to take $p = 2k\lceil \log(2k) \rceil + 1$.

3.2 Ignoring vertex weights

The results of the previous subsection show how to obtain a nonsingular weight matrix from a singular one, at the cost of introducing vertex weights. In this section we show that the vertex weights can be ignored in any proof of #P-hardness. Specifically, we give a polynomial-time reduction from EVAL(A) to EVAL(A, D), for any weight matrix A and any matrix D of positive vertex weights.

Recall the definition of ${\Omega_H}^{(i,j)}(P_r)$, as given in (1). We can analogously define

$${\Omega_H}^{(i,j)}(S_2 T_p P_r) = \{X \in \Omega_H(S_2 T_p P_r) \mid X(u_0) = i,\ \ X(u_r) = j\}\,.$$

Lemma 3.8 *Let $p \geq 1$ be some constant value. Let A be a symmetric matrix and let D be a diagonal matrix of positive vertex weights $\{\lambda_i \mid i \in C\}$. Let Π be the diagonal matrix with (i,i) entry equal to $\sqrt{\lambda_i}$ for all $i \in C$. Then*

$$\sum_{X \in \Omega_H{}^{(i,j)}(S_2 T_p P_r)} w_{A,D}(X) = \sqrt{\lambda_i \lambda_j} B^r{}_{ij}$$

where $B = \Pi \tau_p(ADA)\Pi$.

Proof. We prove the result by induction on r, using the fact that

$$Z_{A,D}(S_2 T_p e) = Z_{\tau_p(ADA),D}(e)$$

for any edge e. Hence the result holds for $r = 1$. Suppose that the result holds for some r such that $r \geq 1$. Then

$$\begin{aligned}
\sum_{X \in \Omega_H{}^{(i,j)}(P_{r+1})} w_{A,D}(X) &= \sum_{\ell \in C} \sum_{Y \in \Omega_H{}^{(i,\ell)}(P_r)} w_{A,D}(Y)\, A_{\ell j} \lambda_j \\
&= \sum_{\ell \in C} \sqrt{\lambda_i \lambda_\ell} B^r{}_{i\ell} A_{\ell j} \lambda_j \\
&= \sqrt{\lambda_i \lambda_j} \sum_{\ell \in C} B^r{}_{i\ell}\, B_{\ell j} \\
&= \sqrt{\lambda_i \lambda_j}\, B^{r+1}{}_{ij},
\end{aligned}$$

as required. □

Using this device, we may prove that the vertex-unweighted problem is at least as easy as any vertex-weighted version, for any weight matrix A.

Theorem 3.2 *Let A be a symmetric matrix with no two linearly dependent columns. There is a polynomial-time reduction from* EVAL(A) *to* EVAL(A, D)*, for any diagonal matrix D of positive vertex weights.*

Proof. Let G be a given instance of EVAL(A), and let D be any diagonal matrix of vertex weights. Define Π to be the diagonal matrix with $\Pi_{ii} = \sqrt{\lambda_i}$ for $i \in C$, as in Lemma 2. From G, form the (multi)graph Γ with edge bipartition $E_\Gamma = E' \cup E''$, as follows. Take each vertex $v \in V$ in turn. Let $d(v)$ denote the degree of v in G. If $d(v) \geq 3$ then replace v by three vertices v_1, v_2, v_3. Let $\{\{v_i, v_j\} \mid 1 \leq i < j \leq 3\} \subseteq E''$, and join each neighbour of v in G to exactly one of v_1, v_2 or v_3, using an edge in E'. This can be done so that $d_\Gamma(v_i) \leq d(v)$, where $d_\Gamma(v_i)$ denotes the degree of v_i in Γ. If v is a vertex of degree 2, replace v by two v_1, v_2. Each neighbour of v in G is joined to exactly one of v_1, v_2 by an edge in E', and the edge $\{v_1, v_2\}$ is placed in E'' with multiplicity two (a double edge). Here ensure that $d_\Gamma(v_i) = 3$ for $i = 1, 2$. Finally, if v is a vertex of degree 1 then add the loop $\{v, v\}$ into E''. Consider the degree of v to be 3 in

Γ. Note that each vertex in V_Γ is the endpoint of exactly two edges in E'' (considering each "end" of the loop added to vertices with degree 1 to be a distinct edge).

Let n be the number of vertices in Γ. Fix p to be the value on the right hand side of (9). Let $S_2T_pS_r{}''\Gamma$ denote the graph obtained from Γ by replacing every edge in E'' by $S_2T_pP_r$. That is, $S_2T_pS_r{}''\Gamma = S_2T_pS_r{}^{(E'')}\Gamma$. Form $S_2T_pS_r{}''\Gamma$ for $1 \le r \le (n+1)^k$. This can be achieved in polynomial time. Let $B = \Pi\tau_p(ADA)\Pi$. For ease of notation, let $w_A{}'(X) = w_A{}^{(E')}(X)$ and let $w_A{}''(X) = w_A{}^{(E'')}(X)$ for all $X : V \mapsto C$. It is not difficult to see that

$$\begin{aligned} Z_{A,D}(S_2T_pS_r{}''\Gamma) &= \sum_{X:V_\Gamma\mapsto C} \widetilde{w}_D(X)\, w_A{}'(X) \times \\ &\qquad \prod_{\{v,w\}\in E''} \sum_{Y\in\Omega_H{}^{(X(v),X(w))}(S_2T_pP_r)} w_{A,D}(Y)/(\lambda_{X(v)}\lambda_{X(w)}) \\ &= \sum_{X:V_\Gamma\mapsto C} \widetilde{w}_D(X)\, w_A{}'(X) \prod_{\{v,w\}\in E''} B^r{}_{X(v)X(w)}/\sqrt{\lambda_{X(v)}\lambda_{X(w)}} \\ &= \sum_{X:V_\Gamma\mapsto C} w_A{}'(X) w_{\sigma_r B}{}''(X). \end{aligned} \tag{10}$$

Here the second equality follows by Lemma 3.8, and the third equality follows as every vertex in Γ is the endpoint of exactly two edges in E''.

The right hand side of (10) is of the form specified in Lemma 3.4, with $U = E''$. Moreover, it is independent of the vertex weights. The matrix $\tau_p(ADA)$ is nonsingular, by Theorem 3.1. Therefore the matrix B is nonsingular, so Lemma 3.4 applies. If we knew the value of $Z_{A,D}(S_2T_pS_r{}''\Gamma)$ for $1 \le r \le (n+1)$, we could calculate the value of $\sum_{X:V_\Gamma\mapsto C} w_A{}'(X)\, w_I{}''(X)$ in polynomial time, by Lemma 3.4. However, this quantity is equal to $Z_A(G)$, by inspection. This completes the polynomial-time reduction. □

The following result is a corollary of Lemma 3.5 and Theorem 3.2. It shows that we may always "collapse" vertices in H with the same set of neighbours, into a single vertex. If the resulting graph gives rise to a #P-complete counting problem, then $\#H$ is #P-complete. This result will be used repeatedly in the proof of Theorem 1.1.

Corollary 3.4 *Let H be a graph, and let H' be obtained from H by replacing all vertices with the same neighbourhood structure by a single vertex (with that neighbourhood structure). If $\#H'$ is #P-complete then $\#H$ is also #P-complete.*

Proof. The graph H' described above is the one used in Lemma 3.5. Let A, A' be the adjacency matrices of H and H', respectively. Then EVAL(A', D) and EVAL(A) are equivalent for some diagonal matrix D of positive vertex weights, as in Lemma 3.5. Moreover, there is a polynomial-time reduction from EVAL(A') to EVAL(A', D), by Theorem 3.2. This completes the proof. □

4 The main proof

In this section, we prove Theorem 1.1. First, we describe those graphs H for which the associated counting problem can be solved in polynomial time. Recall that an isolated vertex with no loops is considered to be a complete bipartite graph.

Lemma 4.1 *Suppose that H is a complete graph with all loops present, or a complete bipartite graph with no loops. Then the counting problem $\#H$ can be solved in polynomial time.*

Proof. If H is an isolated vertex without a loop then $\Omega_H(G)$ is empty unless G is a collection of isolated vertices, in which case $|\Omega_H(G)| = 1$. Suppose that H is the complete graph on k vertices with every loop present, where $k \geq 1$. If G has n vertices then $|\Omega_H(G)| = k^n$. Finally, suppose that H is the complete bipartite graph with vertex bipartition $C_1 \cup C_2$, and with no loops present. If G is not bipartite then $\Omega_H(G)$ is empty. Finally, assume that G is bipartite with vertex bipartition $V_1 \cup V_2$. Suppose that $|C_i| = k_i$ for $i = 1, 2$ and $|V_i| = n_i$ for $i = 1, 2$. Then it is not difficult to see that

$$|\Omega_H(G)| = {k_1}^{n_1} {k_2}^{n_2} + {k_1}^{n_2} {k_2}^{n_1}.$$

This completes the proof. $\square$

The next result shows that a counting problem is #P-complete whenever the counting problem associated with at least one of its connected components is #P-complete. This is a critical result for our proof of Theorem 1.1.

Theorem 4.1 *Suppose that H is a graph with connected components $H_1, \ldots, H_T$. If $\#H_\ell$ is #P-complete for some ℓ such that $1 \leq \ell \leq T$, then $\#H$ is #P-complete.*

Proof. Let A, A_ℓ be the adjacency matrix of H, H_ℓ respectively, for $1 \leq \ell \leq T$. Fix a positive integer r such that ${A_\ell}^r$ has only positive entries, for $1 \leq \ell \leq T$. Note that r can be found in constant time. We show how to perform polynomial-time reductions from EVAL(A_ℓ) to EVAL(A), for $1 \leq \ell \leq T$. This is sufficient, since at least one of the problems EVAL(A_ℓ) is #P-hard, by assumption, and EVAL(A) is clearly in #P.

Let G be a given graph. We wish to calculate the values of $Z_{A_\ell}(G)$ in polynomial time, for $1 \leq \ell \leq T$. For $1 \leq s \leq T$, form the graph G_s with edge bipartion $E' \cup E''$ from G as follows. Let $v \in V$ be an arbitrary vertex of G. Take s copies of G, placing all these edges in E'. Let $\{v_i, v_j\} \in E''$ for $1 \leq i < j \leq s$, where v_i is the copy of v in the ith copy of G. These graphs can be formed from G in polynomial time. Now for $1 \leq s \leq T$ and $1 \leq p \leq s^{2k^2}$, form the graph $(S_r T_p)'' G_s$ by taking the p-thickening of each edge in E'', and then forming the r-stretch of each of these $ps(s-1)/2$ edges. That is, between v_i and v_j we have p copies of the path P_r, for $1 \leq i < j \leq s$. These graphs can be formed from G_s in polynomial time. Let $Z_A(G, c)$ be defined by

$$Z_A(G, c) = |\, \{X \in \Omega_H(G) \mid X(v) = c\} \,|.$$

Then

$$\begin{aligned}
&Z_A((S_rT_p)''G_s) \\
&= \sum_{\ell=1}^{T} \sum_{X:\{v_1,\dots,v_s\}\mapsto C_\ell} \prod_{1\le i\le s} Z_{A_\ell}(G, X(v_i)) \prod_{1\le i<j\le s} \left({A^r}_{X(v_i)X(v_j)}\right)^p && (11)\\
&= \sum_{w\in\mathcal{W}^{(s)}(A)} c_w\, w^p, && (12)
\end{aligned}$$

where $\mathcal{W}^{(s)}(A)$ is defined by

$$\mathcal{W}^{(s)}(A) = \left\{ \prod_{1\le i<j\le s} {A^r}_{X(v_i)X(v_j)} \mid X : \{v_1,\dots,v_s\} \mapsto C_\ell \text{ for some } \ell, 1\le \ell \le T \right\} \setminus \{0\}\,.$$

The set $\mathcal{W}^{(s)}(A)$ can be formed explicitly in polynomial time. Arguing as in (6), the set $\mathcal{W}^{(s)}(A)$ has at most s^{2k^2} distinct elements, all of which are positive. Suppose that we knew the values of $Z_A((S_rT_p)''G_s)$ for $1\le p\le |\mathcal{W}^{(s)}(A)|$. Then, by Lemma 3.2, the values c_w for $w\in\mathcal{W}^{(s)}(A)$ can be found in polynomial time. Adding them, we obtain $f_s = f_s(G) = \sum_{w\in\mathcal{W}^{(s)}(A)} c_w$. This value is also obtained by putting $p=0$ in (12). Equating this to the value obtained by putting $p=0$ in (11), we see that

$$\begin{aligned}
f_s &= \sum_{\ell=1}^{T} \sum_{X:\{v_1,\dots,v_s\}\mapsto C_\ell} \prod_{1\le i\le s} Z_{A_\ell}(G, X(v_i)) \\
&= \sum_{\ell=1}^{T} Z_{A_\ell}(G)^s.
\end{aligned}$$

For ease of notation, let $x_\ell = Z_{A_\ell}(G)$ for $1\le \ell\le T$. We know the values of $f_s = \sum_{\ell=1}^{T} {x_\ell}^s$ for $1\le s\le T$. Let ψ_s be the sth elementary symmetric polynomial in variables $x_1,\dots,x_T$, defined by

$$\psi_s = \sum_{1\le i_1<\cdots<i_s\le T} x_{i_1}\cdots x_{i_s}$$

for $1\le s\le T$. Now

$$f_s - \psi_1\, f_{s-1} + \cdots + (-1)^{s-1} f_1\, \psi_{s-1} + (-1)^s\, s\, \psi_s = 0$$

for $1\le s\le T$ (this is Newton's Theorem, see for example [8, p. 12]). Using these equations, we can evaluate ψ_s for $1\le s\le T$ in polynomial time. But $x_1,\dots,x_T$ are the roots of the polynomial

$$g(z) = z^T - \psi_1 z^{T-1} + \cdots + (-1)^{T-1}\psi_{T-1} z + (-1)^T \psi_T.$$

Since this is a polynomial with integral coefficients, the roots can be found in polynomial time using the algorithm of Lenstra, Lenstra and Lovász [17]. Thus we obtain the set of values $\{Z_{A_\ell}(G) \mid 1\le \ell\le T\}$.

Let $N = |\{Z_{A_\ell}(G) \mid 1 \le \ell \le T\}|$. If $N = 1$ then all the values of $Z_{A_\ell}(G)$ are equal. Thus we know the value of $Z_{A_\ell}(G)$ for $1 \le \ell \le G$, as required. Otherwise, search for a connected graph Γ, with the minimal number of vertices, such that $|\{Z_{A_\ell}(\Gamma) \mid 1 \le \ell \le T\}| = N$. We know that Γ exists, since it is a minimal element of a nonempty set with a lower bound (the empty graph), using partial order on graphs defined by the number of vertices and inclusion. Moreover, Γ depends only on H. Therefore we can find Γ by exhaustive search, in constant time. (This constant may very well be huge, but we are not seeking a practical algorithm.) We also know the values $Z_{A_\ell}(\Gamma)$ for $1 \le \ell \le T$. Let $\sim$ be the equivalence relation on $\{1, \dots, T\}$ such that $Z_{A_\ell}(\Gamma) = Z_{A_s}(\Gamma)$ if and only if $\ell \sim s$. Let π be the partition of $\{1, \dots, T\}$ consisting of the equivalence classes of $\sim$. Write $\pi = I_1 \cup \cdots \cup I_N$ and let $\mu_j = |I_j|$ for $1 \le j \le N$. Finally, let $\mu = \max\{\mu_j \mid 1 \le j \le N\}$. Assume without loss of generality that $j \in I_j$ for $1 \le j \le N$. That is, the first N values of $Z_{A_\ell}(\Gamma)$ form a transversal of the N equivalence classes.

We perform a second reduction, which is an adaption of the one just described. For $1 \le s \le \mu$ and $1 \le t \le N$, form the graph $G_{(s,t)}$ with edge bipartition $E' \cup E''$ as follows. Let w be an arbitrary vertex in Γ, and recall the distinguished vertex v in G. Take s copies of G and t copies of Γ, placing all these edges in E'. Let $V^* = \{w_1, \dots, w_s\} \cup \{v_1, \dots, v_t\}$, where w_i is the copy of w in the ith copy of Γ and v_j is the copy of v in the jth copy of G. Finally, let E'' be the set of all possible edges between the vertices in V^*. Form the graph $(S_rT_p)''G_{(s,t)}$ for $1 \le p \le (s+t)^{2k^2}$ and $1 \le s \le N$, by replacing each edge in E'' by p copies of the path P_r. Arguing as in the first reduction, the values of $Z_A((S_rT_p)''G_{(s,t)})$ for $1 \le p \le (s+t)^{2k^2}$ can be used to produce the values

$$f_{(s,t)}(G) = \sum_{\ell=1}^{T} Z_{A_\ell}(G)^s \, Z_{A_\ell}(\Gamma)^t \tag{13}$$

for $1 \le s \le \mu$, $1 \le t \le N$, in polynomial time.

We can rewrite (13) as

$$f_{(s,t)}(G) = \sum_{j=1}^{N} \left(\sum_{\ell \in I_j} Z_{A_\ell}(G)^s \right) Z_{A_j}(\Gamma)^t.$$

For each fixed value of s, we know the value $f_{(s,t)}(G)$ for $1 \le t \le N$. First suppose that $Z_{A_\ell}(\Gamma) \ne 0$ for $1 \le \ell \le T$. Using Lemma 3.2, we obtain the coefficients ${c^{(s)}}_j = \sum_{\ell \in I_j} Z_{A_\ell}(G)^s$, for $1 \le j \le N$, in polynomial time. We can do this for $1 \le s \le \mu$. Now suppose without loss of generality that $Z_{A_1}(\Gamma) = 0$. Then Lemma 3.2 only guarantees that we can find ${c^{(s)}}_j$ for $2 \le j \le N$, in polynomial time. However, we know the set $\{Z_{A_\ell}(G) \mid 1 \le \ell \le T\}$. Therefore we can form the value $c^{(s)} = \sum_{\ell=1}^{T} Z_{A_\ell}(G)^s$ in polynomial time, for $1 \le s \le \mu$. Then

$${c^{(s)}}_1 = \sum_{\ell \in I_1} Z_{A_\ell}(G)^s = c^{(s)} - \sum_{j=2}^{N} {c^{(s)}}_j.$$

Thus in both cases we can find the values of

$$c^{(s)}{}_j = \sum_{\ell \in I_j} Z_{A_\ell}(G)^s$$

for $1 \le j \le N$ and $1 \le s \le \mu$, in polynomial time. Arguing as above, using Newton's Theorem, we can find the set of values $\{Z_{A_\ell}(G) \mid \ell \in I_j\}$ for $1 \le j \le N$ in polynomial time. If all these values are equal, then we know all the values $Z_{A_\ell}(G)$ for $\ell \in I_j$. Otherwise, we perform the second reduction again, for the graph $H_{I_j} = \cup_{\ell \in I_j} H_\ell$. We obtain a tree of polynomial-time reductions, where each internal node has at least two children, and there are at most T leaves. (A leaf is obtained when all values $Z_{A_\ell}(G)$ in the cell of the partition are equal, which will certainly happen when the cell is a singleton set.) There are at most T internal nodes in such a tree. That is, we must perform at most $T+1$ reductions in all. This guarantees that we can obtain all the values $Z_{A_\ell}(G)$ for $1 \le \ell \le T$ in polynomial time, as required. □

Proof of Theorem 1.1.

The remainder of the section is devoted to the proof of Theomem 1.1. Let H be a graph with adjacency matrix A, and let G be an arbitrary graph. We can assume that G is connected, and by Theorem 4.1 we can also assume that H is connected. By Lemma 4.1 we can assume that H is not a complete graph with all loops present, or a complete bipartite graph with no loops present. The problem $\#H$ is clearly in #P, so it remains to show that it is #P-hard. We do this by demonstrating a series of polynomial-time reductions from some known #P-hard counting problem to $\#H$.

Case 1. Suppose that H has a loop on every vertex. Then H is not complete. Let G be a given connected graph. Form the graph G' from G by introducing a new vertex v_0 and joining all vertices of G to v_0. The graph G' can be formed from G in polynomial time. Define H_i to be the subgraph of H induced by the set of neighbours of the vertex i in H. Note that $i \in H_i$ since H has a loop on every vertex. This implies that H_i is connected for all $i \in C$. Let H' be the graph with connected components given by the multiset $\{H_i \mid i \in C\}$. Then

$$|\Omega_H(G')| = \sum_{i=1}^{k} |\Omega_{H_i}(G)| = |\Omega_{H'}(G)|.$$

This gives a polynomial-time reduction from $\#H'$ to $\#H$. Therefore, by Theorem 4.1, it suffices to show that $\#H_i$ is #P-hard for some $i \in C$. We will iterate this reduction until a #P-hard problem is obtained.

For $i \in C$ let $S_i = \{j \in C \mid \{i,j\} \in E_H, i \ne j\}$. That is, S_i is the set of neighbours of i in H distinct from i. Suppose that there exists a vertex $i \in C$ which satisfies the following conditions:

(i) there exists $j \in C$ such that $\{i,j\} \notin E_C$,

(ii) the subgraph of H induced by S_i is not a clique.

Then H_i is a connected graph which (by (i)) is smaller than H. There is still a loop on every vertex of H_i. By (ii), the graph H_i is not complete. Repeat this process with $H = H_i$. In a finite number of steps we reach a graph H_i such that no vertex of H_i satisfies the given conditions. There is at least one vertex (such as i) attached to all other vertices. The subgraph of H_i induced by S_j is not a clique, for all such vertices j. For all other vertices j, the subgraph of H_i induced by S_j is a clique.

Using Corollary 3.4, we can collapse all vertices of H_i with the same neighbourhood structure down to a single vertex. Suppose that there are now $q+1$ vertices in H_i. We know that $q \geq 2$, since the subgraph of H_i induced by S_i is not a clique. The graph H_i encodes the Widom–Rowlinson model of a gas with q particles. Let A be the adjacency matrix of H_i, and let I_q be the $q \times q$ identity matrix. The matrix A is shown below, together with its square:

$$A = \left[\begin{array}{c|ccc} 1 & 1 & \cdots & 1 \\ \hline 1 & & & \\ \vdots & & I_q & \\ 1 & & & \end{array}\right], \qquad A^2 = \begin{bmatrix} q+1 & 2 & 2 & \cdots & 2 \\ 2 & 2 & 1 & \cdots & 1 \\ 2 & 1 & 2 & \cdots & 1 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 2 & 1 & 1 & \cdots & 2 \end{bmatrix}.$$

There is a polynomial time reduction from EVAL(A^2) to EVAL(A), using Lemma 3.1 with $r = 2$ and $p = 1$. Let B be obtained from A^2 by replacing all entries which do not equal 1 by 0. By Corollary 3.3, there is a polynomial-time reduction from EVAL(B) to EVAL(A^2). However, EVAL(B) is the problem of counting proper q-colourings of a graph. This problem is #P-hard for $q \geq 3$.

Finally suppose that $q = 2$. Then A^2 has distinct entries 3, 2 and 1 which are coprime. Apply Corollary 3.2 with $\mathcal{S} = \{2\}$. This shows that there is a polynomial time reduction from EVAL(B) to EVAL(A^2), where B is the matrix shown below together with its square.

$$B = \begin{bmatrix} 0 & 1 & 1 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix}, \qquad B^2 = \begin{bmatrix} 2 & 1 & 1 \\ 1 & 2 & 1 \\ 1 & 1 & 2 \end{bmatrix}.$$

By Lemma 3.1, there is a polynomial time reduction from EVAL(B^2) to EVAL(B). Now apply Corollary 3.3 to B^2. This gives a polynomial-time reduction to EVAL(B^2), from the #P-hard problem of counting proper 3-colourings. Hence EVAL(A^2) is #P-hard, completing the proof for Case 1.

Case 2. Now suppose that H has some looped vertices and some unlooped vertices. We use the same reduction as Case 1. Recall the subgraph H_i of H, induced by the set of all neighbours of i. Here $i \in H_i$ if and only if there is a loop at i. Since H is connected and contains both looped and unlooped vertices, there is an edge $\{i, j\} \in E_H$ such that i is looped and j is unlooped. Note that the edge $\{i, j\}$, together with the loop at i, describes the #P-hard problem of counting independent sets in graphs. Consider the graph H_i. This is a connected graph which has some looped vertices and some unlooped vertices. In addition, at least one looped vertex of H_i is joined to all the vertices of H_i. Using Corollary 3.4, we can assume that i is the only such vertex.

Now suppose that j and ℓ are both unlooped neighbours of i which are joined by an edge. Then $(H_i)_\ell$ is smaller than H_i, since it does not contain ℓ. It still contains independent sets as a subproblem. Therefore we can replace H_i by $(H_i)_\ell$. After a finite number of steps we can assume that there are no edges between unlooped vertices in H_i.

Next, suppose that there exists more than one looped vertex in H_i. By above, if $\ell \neq i$ and ℓ has a loop, then ℓ is not joined to all the vertices in H_i. If ℓ is a looped vertex which is joined to an unlooped vertex j, the graph $(H_i)_\ell$ is smaller than H_i and still contains independent sets as a subproblem. Hence we can replace H_i by $(H_i)_\ell$. After a finite number of steps, we can assume that i is the only looped vertex which is joined to any unlooped vertices.

Suppose that there exists a looped vertex ℓ in H_i such that the subgraph of H_i generated by S_ℓ is not a clique (where, recall, S_ℓ is the set of all neighbours of S_ℓ in H_i other than ℓ). Then $(H_i)_\ell$ has a loop on every vertex but is not complete. This problem is #P-hard, by Case 1. Otherwise, the subgraph of H_i generated by S_ℓ is a clique, for all looped vertices in H_i other than i. Using Corollary 3.4, we can collapse all vertices with the same neighbourhood structure down to a single vertex. The resulting graph has one unlooped vertex and q looped vertices, where $q \geq 2$. There is one looped vertex i which is joined to all the other vertices. (We can think of this as the graph for q-particle Widom–Rowlinson, with the loop removed from one low-degree vertex.) The adjacency matrix A of this graph is shown below, together with its square:

$$A = \begin{bmatrix} 1 & 1 & 1 & \cdots & 1 & 1 \\ 1 & 1 & 0 & \cdots & 0 & 0 \\ 1 & 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 1 & 0 & 0 & \cdots & 1 & 0 \\ 1 & 0 & 0 & \cdots & 0 & 0 \end{bmatrix}, \qquad A^2 = \begin{bmatrix} q+1 & 2 & 2 & \cdots & 2 & 1 \\ 2 & 2 & 1 & \cdots & 1 & 1 \\ 2 & 1 & 2 & \cdots & 1 & 1 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 2 & 1 & 1 & \cdots & 2 & 1 \\ 1 & 1 & 1 & \cdots & 1 & 1 \end{bmatrix}.$$

Let B be obtained from A^2 by replacing all entries which do not equal 1 by 0. Using Corollary 3.3, there is a polynomial-time reduction from EVAL(B) to EVAL(A^2). But B describes a graph which has exactly one looped vertex which is joined to all the other vertices, of which there are at least two.

We have reached the point where H contains exactly one looped vertex, joined to all other vertices, of which there are at least one. Arguing as above, we can assume that there are no edges between the unlooped vertices. Then we can replace all the unlooped vertices by a single vertex, using Corollary 3.4. The resulting graph describes the #P-hard problem of counting independent sets in graphs. Thus $\#H_i$ is #P-hard , and so is the original problem $\#H$. This completes the proof in Case 2.

Case 3. Now suppose that H is a bipartite graph with no loops. Let G be a given connected graph. If G is not bipartite then $\Omega_H(G) = \emptyset$. Therefore we can assume that G is bipartite, with vertex partition $V_1 \cup V_2$. We adapt the reduction used in Cases 1 and 2 to the bipartite case, by using two apices instead of one. Specifically, form the graph G' from G by introducing two new vertices v_1, v_2, and joining all vertices in V_i to v_i, for $i = 1, 2$. Also let $\{v_1, v_2\}$ be an edge in G'. The graph G' can be formed from G in polynomial time.

Let N_i denote the set of neighbours of i in H, for all $i \in C$. If $\{i, j\} \in E_H$, let H_{ij} be the subgraph of H induced by $N_i \cup N_j$. Let H' be the graph with connected components given by the multiset $\{H_{ij} \mid \{i, j\} \in E_H\}$. Then

$$|\Omega_H(G')| = \sum_{\{i,j\} \in E_H} |\Omega_{H_{ij}}(G)| = |\Omega_{H'}(G)|.$$

This gives a polynomial-time reduction from $\#H'$ to $\#H$. Therefore, by Theorem 4.1, it suffices to show that $\#H_{ij}$ is #P-hard for some $\{i, j\} \in E_H$.

The diameter of a connected graph is the maximum, over all vertices i, j, of the length of the shortest path between i and j. Since H is bipartite and not complete it has diameter at least 3. Suppose that the diameter of H is $d \geq 5$. Let A be the adjacency matrix of H, and consider the matrix A^3. Then A^3 is the weight matrix of a bipartite graph $\tilde{H}$ with diameter strictly between 3 and $d - 2$. There is a polynomial-time reduction from $\#\tilde{H}$ to EVAL(A), by Corollary 3.1. After a finite number of steps, we may assume that H has diameter 3 or 4.

Suppose that H has vertex bipartition $C = C_1 \cup C_2$ where $|C_1| = r$ and $|C_2| = s$. Say that i is on the left if $i \in C_1$, otherwise say that i is on the right. Let $x \in C_1$, $y \in C_2$ be such that $\{x, y\} \notin E_H$. Such a pair exists, since H is not a complete bipartite graph. Since H has diameter 3 or 4, there exist $i \in C_1$ and $j \in C_2$ such that

$$x \mapsto j \mapsto i \mapsto y$$

is a path in H between x and y. Consider the graph H_{ij}. Note that i and j both belong to H_{ij}, and that H_{ij} is a connected bipartite graph with no loops which is not complete. Moreover, all vertices on the right in H_{ij} are joined to i, and all vertices on the left in H_{ij} are joined to j. Using Corollary 3.4, we can assume that i and j are the only vertices in H_{ij} which satisfy these conditions. Let A be the adjacency matrix of H_{ij}. Then

$$A = \begin{bmatrix} 0 & B \\ B^T & 0 \end{bmatrix},$$

where B has the form

$$B = \begin{bmatrix} 1 & 1 & \cdots & 1 \\ 1 & * & \cdots & * \\ \vdots & \vdots & \ddots & \vdots \\ 1 & * & \cdots & * \end{bmatrix}.$$

Here '$*$' stands for entries which we have yet to determine. Note that B is not necessarily square. It is an $r \times s$ matrix where $r \geq 2$ and $s \geq 2$.

If all of the entries marked '$*$' are equal to zero, we are done. Otherwise, choose $\{i', j'\} \in E_{H_{ij}}$ such that $i' \neq i$, $j' \neq j$ and

$$\deg(i') + \deg(j')$$

is maximal (referring to the degree of i' and j' in H_{ij}). The graph $(H_{ij})_{i'j'}$ is smaller than H_{ij}. If it is not a complete bipartite graph, we may work with this graph instead. Using Corollary 3.4, we can collapse all vertices with the same neighbourhood structure into a single vertex. Thus we can assume that B has the form

$$B = \begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & 1 & 0 & \cdots & 0 \\ 1 & 0 & * & \cdots & * \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 0 & * & \cdots & * \end{bmatrix}.$$

If all the entries marked '$*'$ are zero, we are done. Repeat this procedure, which must terminate in a finite number of steps. Finally, apply Corollary 3.4 again, to delete any repeated rows of A. The resulting bipartite graph H has the following form. The vertex bipartition of H is $C_1 \cup C_2$, where $C_1 = \{i_1, \ldots, i_r\}$ and $C_2 = \{j_1, \ldots, j_s\}$ where $s \in \{r, r+1\}$. All edges of the form $\{i_1, j_\ell\}$, $\{i_\ell, j_1\}$, $\{i_\ell, j_\ell\}$ are present for $1 \leq \ell \leq r$, apart from possibly the edge $\{i_r, j_r\}$ in the case that $s = r$. Moreover, $r \geq 2$ unless $s = r$ and the edge $\{i_r, j_r\}$ is present, in which case $r \geq 3$.

Let A be the adjacency matrix of H. The matrix A^2 is a block diagonal matrix with two blocks, one given by BB^T and one given by B^TB. By Theorem 4.1, it suffices to show that the problem associated with at least one of these blocks is #P-hard. Using Corollary 3.3, there is a polynomial-time reduction from EVAL(F) to EVAL(BB^T), where F is obtained from BB^T by replacing all entries which do not equal 1 by zero. When $s = r$ but the edge $\{i_r, j_r\}$ is absent, the graph corresponding to F is connected and has a looped vertex and an unlooped vertex. Therefore EVAL(F) is #P-hard, by Case 2. Otherwise, the matrix F describes the problem of counting proper $(r-1)$-colourings of graphs. This is #P-hard when $r \geq 4$.

Suppose now that $s = r = 3$ and the edge $\{i_3, j_3\}$ is present. Then the distinct entries of B^2 are coprime. $\mathcal{S} = \{2\}$ to give a polynomial-time reduction from the problem EVAL(F) to EVAL(B^2), where

$$F = \begin{bmatrix} 0 & 1 & 1 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix}.$$

The graph corresponding to F is connected and has a looped vertex and an unlooped vertex. Therefore the corresponding counting problem is #P-hard, by Case 2. This completes the proof when $s = r = 3$ and the edge $\{i_3, j_3\}$ is present.

Next suppose that $r = 2$ and $s = 3$. Here, the distinct entries of BB^T are coprime. Therefore we can apply Corollary 3.2 with $\mathcal{S} = \{2\}$. This gives a polynomial-time reduction to EVAL(BB^T) from the problem of counting independent sets in graphs. The latter is #P-hard. Finally, suppose that $r = 3$ and $s = 4$. Then $B^T B$ is given by

$$B^T B = \begin{bmatrix} 3 & 2 & 2 & 1 \\ 2 & 2 & 1 & 1 \\ 2 & 1 & 2 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix}.$$

Let F be the matrix obtained from $B^T B$ by replacing all entries which do not equal 2 by 0, and those entries equal to 2 by 1. Using Corollary 3.2 with $\mathcal{S} = \{2\}$ again, there is a polynomial-time reduction from EVAL(F) to EVAL(BB^T). But F describes a graph which has a connected component consisting of both looped and unlooped vertices. This problem is #P-hard, by Theorem 4.1 and Case 2. This completes the proof for Case 3.

Case 4. Finally, suppose that H has no loops and is not bipartite. Note that the decision problem corresponding to H is NP-complete, by Hell and Nešetřil [11]. (However, this does not immediately imply that the counting problem is #P-hard.) Recall the reduction used in Cases 1 and 2. It suffices to show that H_i is #P-hard for some $i \in C$, where H_i is the graph induced by the set of neighbours of i in H. Note that $i \notin H_i$ as H has no loops. Suppose that some connected component of H_i is not bipartite, for some $i \in H$. (Recall that we are treating isolated vertices as bipartite graphs.) Then we can replace H by H_i, which is smaller than H. After a finite number of steps we can assume that H is a connected, loopless, nonbipartite graph, and that every component of H_i is a complete bipartite graph, for all $i \in C$. Say H satisfies Property 1 when this holds. (We can assume that H_i is complete since we can apply Case 3 otherwise.)

Next, let d be the minimum length of an odd cycle in H. Since H is not bipartite, we know that $d \geq 3$. If $d \geq 5$ then let A be the adjacency matrix of H, and consider the matrix A^3. Then A^3 is the weight matrix of a graph $\tilde{H}$, which contains all the edges of H and also some new edges. Now $\tilde{H}$ is not bipartite since it still contains the cycle C of minimal length. Also $\tilde{H}$ has no loops as $d \geq 5$. Finally, some chords have been introduced between vertices of the cycle C, showing that the minimal length odd cycle in $\tilde{H}$ is strictly between 3 and $d - 2$. Thus after a finite number of steps we can assume that H contains a triangle. If H does not satisfy Property 1 then we can repeat the entire procedure from the beginning. This can only continue for a finite number of steps, since the first stage deletes at least one vertex and the second stage introduces at least one new edge. Thus we can assume that H satisfies Property 1 and contains a triangle.

Now consider the following reduction. Let G be a given graph, which we can assume is connected. Form the 2-stretch $S_2 G$ of G, by subdividing each edge of G. Join all the newly formed vertices (the midpoints of the edges of G) to another new vertex v_0. Denote the resulting graph by G'. Let N_i denote the set of neighbours of i in H, for all $i \in C$. For all $i \in C$ let V_i be the the set of vertices

$$V_i = \bigcup_{\{i,j\} \in E_H} N_j.$$

So V_i is the vertex i, together with all the neighbours of neighbours of i in H. Let $A^{(i)}$ be the symmetric matrix with rows and columns in one-one correspondence with V_i, defined by

$$A^{(i)}{}_{j,\ell} = |N_i \cap N_j \cap N_\ell|.$$

Finally, let $H^*{}_i$ be the graph underlying the matrix $A^{(i)}$. That is, $\{j, \ell\}$ is an edge in $H^*{}_i$ if and only if $A^{(i)}{}_{j\ell} \neq 0$. Then

$$|\Omega_H(G')| = \sum_{i=1}^{k} \sum_{Y \in \Omega_{H^*{}_i}(G)} w_{A^{(i)}}(Y).$$

Let A' be the matrix with k diagonal blocks given by $A^{(1)}, \ldots, A^{(k)}$. The above says that

$$|\Omega_H(G')| = Z_{A'}(G).$$

Thus it suffices to show that EVAL(A') is #P-hard. By Corollary 3.1, it suffices to show that $\#H'$ is #P-hard, where H' is the graph underlying the weight matrix A'. Using Theorem 4.1, it suffices to show that $\#H^*{}_i$ is #P-hard for some $i \in C$.

Recall that H contains a triangle, $\{i, j, \ell\}$ say. Consider $H^*{}_i$. Then

$$A^{(i)}{}_{ii} = |N_i| \geq 2, \;\; A^{(i)}{}_{ij} = |N_i \cap N_j| \geq 1, \;\; A^{(i)}{}_{i\ell} = |N_i \cap N_\ell| \geq 1.$$

This shows that $H^*{}_i$ has a loop at i and that the edges $\{i, j\}$, $\{i, \ell\}$ are both present in $H^*{}_i$. Finally, note that $\{j, \ell\}$ is not an edge in $H^*{}_i$. For otherwise the graph H_i contains a triangle, contradicting the fact that H satisfies Property 1. Thus $H^*{}_i$ has a connected component containing a looped vertex i and missing an edge $\{j, \ell\}$. This shows that $\#H^*{}_i$ is #P-hard, by Cases 1 and 2 Theorem 4.1. This completes the proof for Case 4. Thus the theorem holds. □

5 Bounded degree graphs

Let A be a weight matrix and let D be an invertible diagonal matrix of vertex weights. Let $\Delta \geq 2$ be a constant. Denote by EVAL$^{(\Delta)}(A, D)$ the problem EVAL(A, D), restricted to those instances G with maximum degree at most Δ. We restate Theorem 1.2 in terms of A and D. The proof is an extension of the proof given in [4] for k-colourings.

Theorem 5.1 *Suppose that A is a symmetric matrix, and D is a diagonal matrix of positive vertex weights. Then there exists a polynomial-time reduction from* EVAL(A, D) *to* EVAL$^{(\Delta)}(A, D)$ *for some constant $\Delta \geq 3$.*

Proof. Let $G = (V, E)$ be a graph with arbitrary maximum degree. We form a graph $\Gamma = (V_\Gamma, E_\Gamma)$ from G, such that Γ has maximum degree at most 3. The vertices of Γ are partitioned into two sets $V_\Gamma = V' \cup V''$, and the edges of Γ are partitioned into two sets $E_\Gamma = E' \cup E''$. To form Γ, perform the following operation in turn to each vertex $v \in V$ with degree $d > 3$: replace v by a copy of P_{d-3} on vertices $v_0, \ldots, v_{d-3}$. All these edges go into E''. Join every neighbour of v in G to exactly one of the new vertices, using an edge in E', in such a way that each v_i has degree 3. All other edges go into E'. Let V'' consist of those vertices which are the endpoint of some edge in E'', and let V' hold all other vertices. Finally, let $U \subseteq V''$ be the set of vertices which are the endpoint of exactly one edge in E''. That is, for each v in g with degree $d > 3$, we have $v_0, v_{d-3} \in U$ and $v_1, \ldots, v_{d-4} \in V'' \setminus U$. Let $m = |E''|$.

Let Π be the diagonal matrix with $\Pi_{ii} = \sqrt{\lambda}_i$, and let p be the least positive integer such that $\tau_p(ADA)$ is nonsingular. Let $B = \Pi\tau_p(ADA)\Pi$. For $1 \leq r \leq (m+1)^k$, form $S_2T_pS_r{}''\Gamma$ by replacing each edge in E'' by $S_2T_pP_r$. Then $S_2T_pS_r{}''\Gamma$ has maximum degree $\Delta = 3p$, for $1 \leq r \leq (m+1)^k$. Let

$$\widetilde{w}_D{}'(X) = \prod v \in V' \lambda_{X(v)}.$$

We will evaluate the function $Z_{A,D}(S_2T_pS_r{}''\Gamma)$ as follows. The first factor is the contribution of the edge weights in E'. The second factor is the contribution of all the vertex weights in V_Γ. Then the contribution from the copies of P_r strung between edges in E'' is recorded, with the vertex weights of the endpoints removed (since they've already been counted). We obtain

$$\begin{aligned}
Z_{A,D}\left(S_2T_pS_r{}''\Gamma\right) &= \sum_{X:V_\Gamma\mapsto C} w_A{}'(X)\ \widetilde{w}_D\,(X) \times \\
&\qquad \prod_{\{v,w\}\in E''}\ \sum_{Y\in\Omega_H{}^{(X(v),X(w))}(S_2T_pP_r)} w_{A,D}(Y)/\lambda_{X(v)}\lambda_{X(w)} \\
&= \sum_{X:V_\Gamma\mapsto C} w_A{}'(X)\ \widetilde{w}_D\,(X) \prod_{\{v,w\}\in E''} B^r{}_{X(v)X(w)}/\sqrt{\lambda_{X(v)}\lambda_{X(w)}} \\
&= \sum_{X:V_\Gamma\mapsto C} w_A{}'(X)\ \widetilde{w}_D{}'(X)\ \widetilde{w}_\Pi{}^{(U)}(X) \prod_{\{v,w\}\in E''} B^r{}_{X(v)X(w)} \\
&= \sum_{X:V_\Gamma\mapsto C} w_A{}'(X)\ \widetilde{w}_D{}'(X)\ \widetilde{w}_\Pi{}^{(U)}(X)\, w_{\sigma_r B}{}''(X). \qquad (14)
\end{aligned}$$

The second equality follows from Lemma 3.8, and the third equality follows since every vertex in V'' belongs to at most two edges in E''. Hence the contribution from the weights of vertices in $V''\setminus U$ disappears, and the contribution from vertices in U is given by the matrix Π. The right hand side of (14) is of the form required by Lemma 3.4, and the matrix B is nonsingular, by choice of p. Therefore Lemma 3.4 applies, with $F = E''$. If we knew the values of $Z_{A,D}\left(S_2T_pS_r{}''\Gamma\right)$ for $1 \leq r \leq (m+1)^k$, we could calculate the value of

$$\sum_{X:V_\Gamma\mapsto C} w_A{}'(X)\ \widetilde{w}_D{}'(X)\ \widetilde{w}_\Pi{}^{(U)}(X)\, w_I{}''(X) \qquad (15)$$

in polynomial time, by Lemma 3.4. But $w_I{}''(X)$ is zero unless all edges in E'' have endpoints coloured with the same colour. Let $d(v)$ denote the degree of vertex $v \in V$ in the graph G. The value in (15) is equal to

$$\sum_{X:V_\Gamma \mapsto C} w_A{}'(X)\ \widetilde{w}_D{}'(X) \prod_{\substack{v\in V_G\\ d(v)>\Delta}} \lambda_{X(v)} = \sum_{X:V_G\mapsto C} w_A(X)\,\widetilde{w}_D(X) = Z_{A,D}(G).$$

This completes the polynomial-time reduction. □

In the above proof, $\Delta = 3p$ where p is the least positive integer such that $\tau_p(ADA)$ is nonsingular. When A is nonsingular, we have $\Delta = 3$. We conjecture that Theorem 1.2 always holds for $\Delta = 3$, although in the singular case this still requires proof. Note that The Widom–Rowlinson model and the Beach model both have nonsingular adjacency matrices, so these problems are still #P-hard when restricted to graphs with maximum degree at most 3.

As mentioned in Section 1, the situation is very different if we consider decision problems rather than counting problems. Theorem 1.2 shows that a #P-complete counting problem remains #P-complete when restricted to instances with some constant maximum degree. However, Galluccio, Hell and Nešetřil [9] showed that there exist graphs H with NP-complete decision problems, such that decision is a polynomial-time operation when restricted to graphs with maximum degree 3. In fact, H can be taken to be a triangle-free graph with chromatic number 3 (see [6] for details).

Theorem 1.2 only concerns problems which are restricted to maximum degree at most Δ, for some constant $\Delta \geq 3$. We show below that $\mathrm{EVAL}^{(2)}(A)$ is always in P, for all weight matrices A. The proof involves the path P_r on $r+1$ vertices and r edges, as defined in the previous section. It also involves the cycle L_r, which has r vertices $u_1, \ldots, u_r$ and r edges $\{u_1, u_r\} \cup \{u_i, u_{i+1} \mid 1 \leq i < r\}$.

Lemma 5.1 *Let A be any symmetric matrix, and let D be a nonsingular diagonal matrix. Then* $\mathrm{EVAL}^{(2)}(A, D) \in \mathrm{P}$. *Let Π the diagonal matrix defined by $\Pi_{ii} = \sqrt{D_{ii}} = \sqrt{\lambda_i}$. Then*

$$Z_{A,D}(P_r) = (D(AD)^r)_{ij} \ \text{ and } \ Z_{A,D}(L_r) = \sum_{i\in C} B^r{}_{ii} = \mathrm{Tr}(B^r).$$

Proof. Let G be any graph with maximum degree at most 2. Then G is a disjoint union of isolated vertices, paths and cycles. Each isolated vertex contributes $\sum_{i\in C} \lambda_i$. We now derive an expression for $Z_{A,D}(P_r)$. Recall the sets $\Omega_H{}^{(i,j)}(P_r)$, defined in (1) for $i, j \in C$, which form a partition of $\Omega_H(P_r)$. Then using Lemma 2 we obtain

$$\begin{aligned}
Z_{A,D}(P_r) &= \sum_{X\in\Omega_H(P_r)} w_{A,D}(X) \\
&= \sum_{i,j\in C} \sum_{X\in\Omega_H{}^{(i,j)}(P_r)} w_{A,D}(X) \\
&= \sum_{i,j\in C} \sqrt{\lambda_i\lambda_j} B^r{}_{ij}. \\
&= (D(AD)^r)_{ij},
\end{aligned}$$

as stated.

Next we derive an expression for $Z_{A,D}(L_r)$. Consider "cutting" the cycle at vertex u_r, replacing it with two vertices u_0 and $u_r{}'$ with edges $\{u_0, u_1\}$ and $\{u_{r-1}, u_r{}'\}$. This gives a bijection between $\Omega_H(L_r)$ and the union over all colours $i \in C$ of $\Omega_H{}^{(i,i)}(P_r)$. Hence

$$Z_{A,D}(L_r) = \sum_{X \in \Omega_H(L_r)} w_{A,D}(X) = \sum_{i \in C} \sum_{X \in \Omega_H{}^{(i,i)}(P_r)} w_{A,D}(X)/\lambda_i = \sum_{i \in C} B^r{}_{ii} = \mathrm{Tr}(B^r),$$

as stated. We divide by λ_i in the second equality since u_1 and $u_r{}'$ are identified in L_r.

Suppose that G has n vertices. We can form the matrices $\{B^r \mid 1 \leq r \leq n\}$ in polynomial time, and use these matrices to calculate $Z_{A,D}(\cdot)$ for each connected component of G. Multiplying these values together, we obtain $Z_{A,D}(G)$ in polynomial time. $\square$

References

[1] M. Albertson, P. Catlin, and L. Gibbons, *Homomorphisms of 3-chromatic graphs, II*, Congressus Numerantium, 47 (1985), pp. 19–28.

[2] G. Bloom and S. Burr, *On unavoidable digraphs in orientations of graphs*, Journal of Graph Theory, 11 (1987), pp. 453–462.

[3] G. R. Brightwell and P. Winkler, *Graph homomorphisms and phase transitions*, Journal of Combinatorial Theory, Series B, (To appear).

[4] R. Bubley, M. Dyer, C. Greenhill, and M. Jerrum, *On approximately counting colourings of small degree graphs*, SIAM Journal on Computing, 29 (1999), pp. 387–400.

[5] R. Burton and J. Steif, *Nonuniqueness of measures of maximal entropy for subshifts of finite type*, Ergodic Theory and Dynamical Systems, 14, pp. 213–236.

[6] P. A. Dreyer, C. Malon, and J. Nešetřil, *Universal H-colorable graphs without a given configuration*, (1999), (Preprint).

[7] M. Dyer and C. Greenhill, *On Markov chains for independent sets*, (1997), (Preprint).

[8] H. M. Edwards, *Galois Theory*, Springer–Verlag, New York, (1984).

[9] A. Galluccio, P. Hell, and J. Nešetřil, *The complexity of H-coloring of bounded degree graphs*, Discrete Mathematics, To appear.

[10] C. Greenhill, *The complexity of counting colourings and independent sets in sparse graphs and hypergraphs*, Computational Complexity, (1999), (To appear).

[11] P. Hell and J. Nešetřil, *On the complexity of H-coloring*, Journal of Combinatorial Theory, Series B, 48 (1990), pp. 92–110.

[12] H. B. Hunt III, M. V. Marathe, V. Radhakrishnan, and R. E. Stearns, *The complexity of planar counting problems*, SIAM Journal on Computing, 27 (1998), pp. 1142–1167.

[13] R. W. Irving, *NP-completeness of a family of graph-colouring problems*, Discrete Applied Mathematics, 5 (1983), pp. 111–117.

[14] F. Jaeger, D. L. Vertigan, and D. Welsh, *On the computational complexity of the Jones and Tutte polynomials*, Mathematical Proceedings of the Cambridge Philosophical Society, 108 (1990), pp. 35–53.

[15] J. Kahn, *An entropy approach to the hard-core model on bipartite graphs*, (1999), (Preprint).

[16] J. L. Lebowitz and G. Gallavotti, *Phase transitions in binary lattice gases*, Journal of Mathematical Physics, 12 (1971), pp. 1129–1133.

[17] A. K. Lenstra, H. W. Lenstra Jnr., and L. Lovász, *Factoring polynomials with rational coefficients*, Mathematische Annalen, 261 (1982), pp. 515–534.

[18] L. A. Levin, *Universal sequential search problems*, Problems of Information Transmission, 9 (1973), pp. 265–266.

[19] N. Linial, *Hard enumeration problems in geometry and combinatorics*, SIAM Journal on Algebraic and Discrete Methods, 7 (1986), pp. 331–335.

[20] M. Luby and E. Vigoda, *Approximately counting up to four*, in *Twenty-Ninth Annual Symposium on Theory of Computing*, ACM, New York, 1997 pp. 682–687.

[21] H. A. Maurier, J. H. Sudborough, and E. Welzl, *On the complexity of the general colouring problem*, Information and Control, 51 (1981), pp. 123–145.

[22] S. P. Vadhan, *The complexity of counting in sparse, regular, and planar graphs*, Preprint (Available from http://www-math.mit.edu/∼salil/), (May 1997).

[23] L. G. Valiant, *The complexity of enumeration and reliability problems*, SIAM Journal on Computing, 8 (1979), pp. 410–421.

[24] D. J. A. Welsh, *Complexity: knots, colourings and counting*, vol. 186 of *London Mathematical Society Lecture Note Series*, Cambridge University Press, Cambridge, (1993).

[25] B. Widom and J. S. Rowlinson, *New model for the study of liquid-vapour phase transition*, The Journal of Chemical Physics, 52 (1970), pp. 1670–1684.