

REFLECTION POSITIVITY, RANK CONNECTIVITY, AND HOMOMORPHISM OF GRAPHS

MICHAEL FREEDMAN, LÁSZLÓ LOVÁSZ, AND ALEXANDER SCHRIJVER

1. INTRODUCTION

For two finite graphs G and H , let $\text{hom}(G, H)$ denote the number of homomorphisms (adjacency-preserving mappings) from G to H . Many interesting graph parameters can be expressed in terms of the number of homomorphisms into a fixed graph: for example, the number of colorings with a number of colors is the number of homomorphisms into a complete graph. Further examples for the occurrence of these numbers in graph theory will be discussed in Section 3.

Another source of important examples is statistical physics, where partition functions of various models can be expressed as graph homomorphism functions. For example, let G be an $n \times n$ grid, and suppose that every node of G (every “site”) can be in one of two states, “UP” or “DOWN”. The properties of the system are such that no two adjacent sites can be “UP”. A “configuration” is a valid assignment of states to each node. The number of configurations is the number of independent sets of nodes in G , which in turn can be expressed as the number of homomorphisms of G into the graph H consisting of two nodes, “UP” and “DOWN”, connected by an edge, and with an additional loop at “DOWN”. To capture more interesting physical models, so-called “vertex coloring models”, one needs to extend the notion of graph homomorphism to the case when the nodes and edges of H have weights (see Section 2.1).

Which graph parameters can be represented as homomorphism functions into weighted graphs? This question is motivated, among others, by the problem of physical realizability of certain graph parameters. Two necessary conditions are easy to prove.

(a) The interaction between two parts of a graph separated by k nodes is bounded by a simple exponential function of k (this will be formalized as *rank-connectivity* in Section 2.2).

(b) Another necessary condition, which comes from statistical mechanics, as well as from extremal graph theory, is *reflection positivity*. Informally, this means that if a system has a 2-fold symmetry, then its partition function is positive. We will formulate a version of this in Section 2.2.

The main result of this paper is to prove that these two necessary conditions are also sufficient (Theorem 2.4). The proof makes use of a simple kind of (commutative, finite-dimensional) C^* -algebras.

Received by the editors July 28, 2004.

2000 *Mathematics Subject Classification*. Primary 05C99; Secondary 82B99.

Key words and phrases. Graph homomorphism, partition function, connection matrix.

This result is related to characterizations of semidefinite functions on commutative semigroups. Let S be a commutative semigroup with identity. A *character* on S is a real-valued multiplicative function; one is interested in describing the convex hull of characters. Let $f : S \rightarrow \mathbf{R}$ be a function on S that is in the convex hull of characters. Then the (infinite) $S \times S$ matrix M_f defined by $(M_f)_{s,t} = f(st)$ is semidefinite. Under various additional conditions on f , this necessary condition is also sufficient; see [1, 2, 3, 6]. Among others, if M_f is positive semidefinite and has finite rank, then f is a convex combination of a finite number of characters.

Our main result can be thought of as treating a specific semigroup (the semigroup \mathcal{PLG} of partially labeled graphs) and functions f on it that are invariant under graph isomorphism (see Section 4.1); on the other hand, instead of a rank condition on the whole matrix M_f , we consider a weaker condition bounding the rank of certain submatrices only. Accordingly, we need infinitely many characters, but these characters will have a finite explicit description as homomorphism numbers.

2. HOMOMORPHISMS, RANK-CONNECTIVITY AND REFLECTION POSITIVITY

2.1. Weighted graph homomorphisms. A *weighted graph* H is a graph with a positive real weight $\alpha_H(i)$ associated with each node i and a real weight $\beta_H(i, j)$ associated with each edge ij .

Let G be an unweighted graph (possibly with multiple edges, but no loops) and H , a weighted graph. To every homomorphism $\phi : V(G) \rightarrow V(H)$, we assign the weights

$$\alpha_\phi = \prod_{u \in V(G)} \alpha_H(\phi(u))$$

and

$$(1) \quad \text{hom}_\phi(G, H) = \prod_{uv \in E(G)} \beta_H(\phi(u), \phi(v))$$

and define

$$(2) \quad \text{hom}(G, H) = \sum_{\substack{\phi: V(G) \rightarrow V(H) \\ \text{homomorphism}}} \alpha_\phi \text{hom}_\phi(G, H).$$

If all the node-weights and edge-weights in H are 1, then this is the number of homomorphisms from G into H (with no weights).

For the purpose of this paper, it will be convenient to assume that H is a complete graph with a loop at all nodes (the missing edges can be added with weight 0). Then the weighted graph H is completely described by a positive integer $d = |V(H)|$, the positive real vector $a = (\alpha_1, \dots, \alpha_d) \in \mathbf{R}^d$ and the real symmetric matrix $B = (\beta_{ij}) \in \mathbf{R}^{d \times d}$. The graph parameter $\text{hom}(\cdot, H)$ will be denoted by f_H .

2.2. Connection matrices of a graph parameter. A *graph parameter* is a function on finite graphs (invariant under graph isomorphism). We allow multiple edges in our graphs, but no loops. A graph parameter f is called *multiplicative* if for the disjoint union $G_1 \cup G_2$ of two graphs G_1 and G_2 , we have $f(G_1 \cup G_2) = f(G_1)f(G_2)$.

A *k-labeled graph* ($k \geq 0$) is a finite graph in which k nodes are labeled by $1, 2, \dots, k$ (the graph can have any number of unlabeled nodes). Two k -labeled graphs are *isomorphic* if there is a label-preserving isomorphism between them.

We denote by K_k the k -labeled complete graph on k nodes, and by O_k , the k -labeled graph on k nodes with no edges. In particular, $K_0 = O_0$ is the graph with no nodes and no edges.

Let G_1 and G_2 be two k -labeled graphs. Their *product* G_1G_2 is defined as follows: we take their disjoint union and then identify nodes with the same label. Hence for two 0-labeled graphs, $G_1G_2 = G_1 \cup G_2$ (disjoint union).

Now we come to the construction that is crucial for our treatment. Let f be any graph parameter. For every integer $k \geq 0$, we define the following (infinite) matrix $M(f, k)$. The rows and columns are indexed by isomorphism types of k -labeled graphs. The entry in the intersection of the row corresponding to G_1 and the column corresponding to G_2 is $f(G_1G_2)$. We call the matrices $M(f, k)$ the *connection matrices* of the graph parameter f .

We will be concerned with two properties of connection matrices, namely their rank and positive semidefiniteness. The rank $r(f, k) = \text{rk}(M(f, k))$, as a function of k , will be called the *rank connectivity function* of the parameter f . This may be infinite, but for many interesting parameters it is finite, and its growth rate will be important for us. We say that a set of k -labeled graphs *spans* if the corresponding rows of $M(f, k)$ have rank $r(f, k)$.

We say that a graph parameter f is *reflection positive* if $M(f, k)$ is positive semidefinite for every k . This property is closely related to the reflection positivity property of certain statistical physical models. Indeed, it implies that for any k -labeled graph G , $f(GG) \geq 0$ (since $f(GG)$ is a diagonal entry of a positive semidefinite matrix). Here the second copy of G can be thought of as a “reflection” (in the set of labeled nodes) of the first.

The condition that $f(GG) \geq 0$ is weaker than the condition that $M(f, k)$ is positive semidefinite, but we can strengthen it as follows. A *quantum graph* is a formal linear combination of a finite number of finite graphs; a *k -labeled quantum graph* is a formal linear combination of a finite number of finite k -labeled graphs. Let $X = \sum_{i=1}^N x_i G_i$ be a k -labeled quantum graph. Then $X^2 = \sum_{G_1, G_2} x_{G_1} x_{G_2} (G_1G_2)$ is a quantum graph that can be obtained by gluing together two copies of X along the labeled nodes. Now if we extend f linearly over quantum graphs, then

$$f(X^2) = x^\top M(f, k)x,$$

showing that the nonnegativity of f over symmetric quantum graphs is equivalent to reflection positivity as we defined it.

2.3. Simple properties of connection matrices. We give a couple of simple facts about connection matrices of a general graph parameter.

Proposition 2.1. *Let f be a graph parameter that is not identically 0. Then f is multiplicative if and only if $M(f, 0)$ is positive semidefinite, $f(K_0) = 1$, and $r(f, 0) = 1$.*

Proof. If f is multiplicative, then $f(K_0)^2 = f(K_0)$, showing that $f(K_0) \in \{0, 1\}$. If $f(K_0) = 0$, then the relation $f(G) = f(GK_0) = f(G)f(K_0)$ implies that $f(G) = 0$ for every G , which was excluded. So $f(K_0) = 1$. Furthermore, $f(G_1G_2) = f(G_1)f(G_2)$ for any two 0-labeled graphs G_1 and G_2 , which implies that $M(f, 0)$ has rank 1 and is positive semidefinite.

Conversely, since $M(f, 0)$ is symmetric, the assumption that $r(f, 0) = 1$ implies that there is a graph parameter ϕ and a constant c such that $f(G_1G_2) =$

$c\phi(G_1)\phi(G_2)$. Since $M(f, 0)$ is positive semidefinite, we have $c > 0$, and so we can normalize ϕ to make $c = 1$. Then $f(K_0) = f(K_0K_0) = \phi(K_0)^2$, whence $\phi(K_0) \in \{-1, 1\}$. We can replace ϕ by $-\phi$, so we may assume that $\phi(K_0) = 1$. Then $f(G) = f(GK_0) = \phi(G)\phi(K_0) = \phi(G)$ for every G , which shows that f is multiplicative. \square

Proposition 2.2. *Let f be a multiplicative graph parameter and $k, l \geq 0$. Then*

$$r(f, k + l) \geq r(f, k) \cdot r(f, l).$$

We will see (see Claim 4.7 below) that in the case when f is reflection positive, the sequence $r(f, k)$ is log-convex. We don't know if this property holds in general.

Proof. Let us call a $(k+l)$ -labeled graph *separated* if every component of it contains either only nodes with label at most k , or only nodes with label larger than k . Consider the submatrix of $M(f, k+l)$ formed by the separated rows and columns. By multiplicativity, this submatrix is the Kronecker (tensor) product of $M(f, k)$ and $M(f, l)$, so its rank is $r(f, k) \cdot r(f, l)$. \square

2.4. Connection matrices of homomorphisms. Fix a weighted graph $H = (a, B)$. For every positive integer k , let $[k] = \{1, \dots, k\}$. For any k -labeled graph G and mapping $\phi : [k] \rightarrow V(H)$, let

$$(3) \quad \text{hom}_\phi(G, H) = \sum_{\substack{\psi: V(G) \rightarrow V(H) \\ \psi \text{ extends } \phi}} \frac{\alpha_\psi}{\alpha_\phi} \text{hom}_\psi(G, H).$$

So

$$(4) \quad \text{hom}(G, H) = \sum_{\phi: [k] \rightarrow V(H)} \alpha_\phi \text{hom}_\phi(G, H).$$

Lemma 2.3. *The graph parameter $f_H = \text{hom}(\cdot, H)$ is reflection positive and $r(f_H, k) \leq |V(H)|^k$.*

Proof. For any two k -labeled graphs G_1 and G_2 and $\phi : [k] \rightarrow V(H)$,

$$(5) \quad \text{hom}_\phi(G_1G_2, H) = \text{hom}_\phi(G_1, H) \text{hom}_\phi(G_2, H).$$

The decomposition (4) writes the matrix $M(f, k)$ as the sum of $|V(H)|^k$ matrices, one for each mapping $\phi : [k] \rightarrow V(H)$; (5) shows that these matrices are positive semidefinite and have rank 1. \square

The main result of this paper is a converse to Theorem 2.3.

Theorem 2.4. *Let f be a reflection positive graph parameter for which there exists a positive integer q such that $r(f, k) \leq q^k$ for every $k \geq 0$. Then there exists a weighted graph H with $|V(H)| \leq q$ such that $f = f_H$.*

3. EXAMPLES

We start with an example showing that exponential growth of rank connectivity is not sufficient in itself to guarantee that a graph parameter is a homomorphism function.

Example 3.1 (Matchings). Let $\Phi(G)$ denote the number of perfect matchings in the graph G . It is trivial that $\Phi(G)$ is multiplicative. We claim that its node rank connectivity is exponentially bounded:

$$r(\Phi, k) = 2^k.$$

Let G be a k -labeled graph, let $X \subseteq [k]$, and let $\Phi(G, X)$ denote the number of matchings in G that match all the unlabeled nodes and the nodes with label in X , but not any of the other labeled nodes. Then we have for any two k -labeled graphs G_1, G_2 ,

$$\Phi(G_1 G_2) = \sum_{X_1 \cap X_2 = \emptyset, X_1 \cup X_2 = [k]} \Phi(G_1, X_1) \Phi(G_2, X_2).$$

This can be read as follows. The matrix $M(\Phi, k)$ can be written as a product $N^T W N$, where N has infinitely many rows indexed by k -labeled graphs, but only 2^k columns, indexed by subsets of $[k]$,

$$N_{G, X} = \Phi(G, X),$$

and W is a symmetric $2^k \times 2^k$ matrix, where

$$W_{X_1, X_2} = \begin{cases} 1 & \text{if } X_1 = [k] \setminus X_2, \\ 0 & \text{otherwise.} \end{cases}$$

Hence the rank of $M(\Phi, k)$ is at most 2^k (it is not hard to see that in fact equality holds).

On the other hand, let us consider K_1 and K_2 as 1-labeled graphs. Then the submatrix of $M(\Phi, 1)$ indexed by K_1 and K_2 is

$$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix},$$

which is not positive semidefinite. Thus $\Phi(G)$ cannot be represented as a homomorphism function.

Example 3.2 (Chromatic polynomial). The following example uses results from [4]. Let $p(G) = p(G, x)$ denote the chromatic polynomial of the graph G . For every fixed x , this is a multiplicative graph parameter. To describe its rank connectivity, we need the following notation. For $k, q \in \mathbf{Z}_+$, let B_{kq} denote the number of partitions of a k -element set into at most q parts. So $B_k = B_{kk}$ is the k -th Bell number. With this notation,

$$r(p, k) = \begin{cases} B_{kx} & \text{if } x \text{ is a nonnegative integer,} \\ B_k & \text{otherwise.} \end{cases}$$

Note that this is always finite, but if $x \notin \mathbf{Z}_+$, then it grows faster than c^k for every c .

Furthermore, $M(p, k)$ is positive semidefinite if and only if either x is a positive integer or $k \leq x + 1$. The parameter $M(p, k)$ is reflection positive if and only if this holds for every k , i.e., if and only if x is a nonnegative integer, in which case indeed $p(G, x) = \text{hom}(G, K_x)$.

The following two examples show that there are important graph parameters that are not defined as homomorphism functions, but that can also be represented as homomorphism functions in a nontrivial way. In fact, it is easier to check that

the conditions in Theorem 2.4 hold than to show that these graph parameters are homomorphism functions, and they first came up as counterexample candidates.

Example 3.3 (Tutte polynomial). The Tutte polynomial is a 2-variable generalization of the chromatic polynomial. It has been observed (see e.g. [9]) that this polynomial behaves similarly along certain hyperbolas. This is also the case from the point of view of homomorphism functions.

We in fact consider the following version $\hat{T}(x, y)$ of the Tutte polynomial (which only differs from the usual version $T(x, y)$ by scaling). For every graph $G = (V, E)$ on n nodes and real parameters x, y , let

$$\hat{T}(G; x, y) = (1-x)^{c(E)}(1-y)^n T(G; x, y) = \sum_{A \subseteq E(G)} (1-x)^{c(A)}(1-y)^{|A|+c(A)},$$

where for $A \subseteq E$, $c(A)$ denotes the number of connected components of the graph (V, A) .

Through this scaling we lose the covariance under matroid duality; but we gain that the following “skein relation” holds for all edges e :

$$(6) \quad T(G; x, y) = (1-y)T(G \setminus e; x, y) - T(G/e; x, y).$$

The chromatic polynomial is $\hat{T}(1-q, 0)$. One can show that the Tutte polynomial \hat{T} behaves just like the chromatic polynomial. Let us assume that $y \neq 1$ (else $T(G; x, y) = 0$ for all nonempty graphs). If $(x-1)(y-1)$ is a positive integer, then $\hat{T}(G; x, y)$ is reflection positive and has exponential rank connectivity; else $\hat{T}(G; x, y)$ is not reflection positive and its rank connectivity is B_k , which is superexponential.

Thus for the case when $(x-1)(y-1)$ is a positive integer, $\hat{T}(G; x, y)$ can be represented as $\text{hom}(G, H)$ for some H . However, except for the case of the chromatic polynomial, and the flow polynomial to be discussed in the next example, we don't know an explicit construction for H .

Example 3.4 (Flows). Let us start with a simple special case. Let $f(G) = 1$ if G is eulerian (i.e., all nodes have even degree), and $f(G) = 0$ otherwise. To represent this function as a homomorphism function, let

$$a = \begin{pmatrix} 1/2 \\ 1/2 \end{pmatrix}, \quad B = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}.$$

It was noted by de la Harpe and Jones [5] that for the weighted graph $H = (a, B)$ we have $\text{hom}(G, H) = f(G)$.

It follows that for this function, reflection positivity holds, and the rank connectivity is at most 2^k .

This example can be generalized quite a bit. Let Γ be a finite abelian group and let $S \subseteq \Gamma$ be such that S is closed under inversion. For any graph G , fix an orientation of the edges. An S -flow is an assignment of an element of S to each edge such that for each node v , the product of elements assigned to edges entering v is the same as the product of elements assigned to the edges leaving v . Let $f(G)$ be the number of S -flows. This number is independent of the orientation.

The choice $\Gamma = \mathbf{Z}_2$ and $S = \mathbf{Z}_2 \setminus \{0\}$ gives the special case above (incidence function of eulerian graphs). If $\Gamma = S = \mathbf{Z}_2$, then $f(G)$ is the number of eulerian subgraphs of G . Perhaps the most interesting special case is when $|\Gamma| = t$ and

$S = \Gamma \setminus \{0\}$, which gives the number of nowhere zero t -flows, also known as the flow polynomial, which can be expressed as $\hat{T}(0, 1 - q)$.

Surprisingly, the parameter $f(G)$ can be described as a homomorphism function. Let Γ^* be the character group of Γ . Let H be the complete directed graph (with all loops) on Γ^* . Let $\alpha_\chi := 1/|\Gamma|$ for each $\chi \in \Gamma^*$, and let

$$\beta_{\chi, \chi'} := \sum_{s \in S} \chi^{-1}(s) \chi'(s),$$

for $\chi, \chi' \in \Gamma^*$.

We show that

$$(7) \quad f = \text{hom}(\cdot, H).$$

Let $n = |V(G)|$ and $m = |\Gamma|$. For any coloring $\phi : V(G) \rightarrow S$ and node $v \in V(G)$, let

$$\partial_\phi(v) = \sum_{\substack{u \in V(G) \\ uv \in E(G)}} \phi(uv) - \sum_{\substack{u \in V(G) \\ vu \in E(G)}} \phi(vu).$$

So ϕ is an S -flow if and only if $\partial_\phi = 0$. Consider the expression

$$A = \sum_{\phi: E(G) \rightarrow S} \prod_{v \in V(G)} \sum_{\chi \in \Gamma^*} \chi(\partial_\phi(v)).$$

The summation over χ is 0 unless $\partial_\phi(v) = 0$, in which case it is m . So the product over $v \in V(G)$ is 0 unless ϕ is an S -flow, in which case it is m^n . So $A \cdot m^{-n}$ counts S -flows.

On the other hand, we can expand the product over $v \in V(G)$; each term will correspond to a choice of a character ψ_v for each v , and so we get

$$A = \sum_{\phi: E(G) \rightarrow S} \sum_{\psi: V(G) \rightarrow \Gamma^*} \prod_{v \in V(G)} \psi_v(\partial_\phi(v)).$$

Here (using that ψ_v is a character)

$$\psi_v(\partial_\phi(v)) = \prod_{\substack{u \in V(G) \\ uv \in E(G)}} \psi_v(\phi(uv)) \prod_{\substack{u \in V(G) \\ vu \in E(G)}} \psi_v(\phi(vu))^{-1},$$

so we get that

$$A = \sum_{\phi: E(G) \rightarrow S} \sum_{\psi: V(G) \rightarrow \Gamma^*} \prod_{uv \in E(G)} \psi_v(\phi(uv)) \psi_u(\phi(uv))^{-1}.$$

Interchanging the summation, the inner sum factors:

$$\begin{aligned} & \sum_{\phi: E(G) \rightarrow S} \prod_{uv \in E(G)} \psi_v(\phi(uv)) \psi_u(\phi(uv))^{-1} \\ &= \prod_{uv \in E(G)} \sum_{s \in S} \psi_v(s) \psi_u(s)^{-1} \\ &= \prod_{uv \in E(G)} \beta_{\psi_u \psi_v} = m^n \text{hom}_\psi(G, H), \end{aligned}$$

showing that

$$A = m^n \text{hom}(G, H).$$

This proves (7).

Example 3.5 (The role of multiple edges). Let us give an example of a reflection positive graph parameter f for which $r(f, k)$ is finite for every k but has superexponential growth. The example also shows that we have to be careful with multiple edges. Let, for each graph G , G' denote the graph which we obtain from G by keeping only one copy of each parallel class of edges. Let

$$f(G) = 2^{-|E(G')|}.$$

It is not hard to see that the connection matrices $M(f, k)$ are positive semidefinite. This graph parameter is, in fact, the limit of parameters of the form $\text{hom}(\cdot, H)$: take homomorphisms into a random graph $H = G(n, 1/2)$, with all node weights $= 1/n$ and all edge weights 1. Furthermore, it is not hard to see that the rank of $M(f, k)$ is $2^{\binom{k}{2}}$. This is finite but superexponential, so the parameter is not of the form $\text{hom}(\cdot, H)$.

Note, however, that for a simple graph G (i.e., if G has no multiple edges), $f(G) = 2^{|E(G)|}$ can be represented as the number of homomorphisms into the graph consisting of a single node with a loop, where the node has weight 1 and the loop has weight $1/2$.

Example 3.6 (Homomorphisms into infinite graphs). We can extend the definition of $\text{hom}(G, H)$ to infinite weighted graphs H provided the node and edge weights form sufficiently fast convergent sequences. Reflection positivity remains valid, but the rank of $M(f_H, k)$ will become infinite, so this graph parameter cannot be represented by a finite H .

More generally, let $a > 0$, $I = [0, a]$, and let $W : I \times I \rightarrow \mathbf{R}$ be a measurable function such that for every $n \in \mathbf{Z}_+$,

$$\int_0^a \int_0^a |W(x, y)|^n dx dy < \infty.$$

Then we can define a graph parameter f_W as follows. Let G be a finite graph on n nodes. Then

$$f_W(G) = \int_{I^n} \prod_{ij \in E(G)} W(x_i, x_j) dx_1 \dots dx_n.$$

It is easy to see that for every weighted (finite or infinite) graph H , the graph parameter f_H is a special case. Furthermore, f_W is reflection positive.

However, it can be shown that the graph parameter in Example 3.5 cannot be represented in this form. A characterization of graph parameters representable in this more general form is given in [7].

4. PROOF OF THEOREM 2.4

4.1. The algebra of graphs. In this first part of the proof, we only assume that for every k , $M(f, k)$ is positive semidefinite, has finite rank r_k , and $r_0 = 1$. We know that f is multiplicative, and hence $f(K_0) = 1$. We can replace the parameter $f(G)$ by $f(G)/f(K_1)^{-|V(G)|}$; this can be reversed by scaling of the node weights of the target graph H representing f , once we have it constructed. So we may assume that $f(K_1) = 1$. Combined with multiplicativity, this implies that we can delete (or add) isolated nodes from any graph G without changing $f(G)$.

It will be convenient to put all k -labeled graphs into a single structure as follows. By a *partially labeled graph* we mean a finite graph in which some of the nodes are labeled by distinct nonnegative integers. Two partially labeled graphs are

isomorphic if there is an isomorphism between them preserving all labels. For two partially labeled graphs G_1 and G_2 , let G_1G_2 denote the partially labeled graph obtained by taking the disjoint union of G_1 and G_2 , and identifying nodes with the same label. This way we obtain a commutative semigroup \mathcal{PLG} . For every finite set $S \subseteq \mathbf{Z}_+$, we call a partially labeled graph S -*labeled*, if its labels form the set S .

It may be useful at this point to discuss the case we are aiming at, namely when $f = \text{hom}(\cdot, H)$ for some weighted graph H . Our assumption that $f(K_1) = 1$ means then that the sum of node weights is 1. For every map $\phi: \mathbf{Z}_+ \rightarrow V(H)$, and every partially labeled graph G with label set S , we define $\text{hom}_\phi(G, H) = \text{hom}_{\phi|_S}(G, H)$. Equation (5) implies that the function $\text{hom}_\phi(\cdot, H)$ is a character of the semigroup \mathcal{PLG} , and formula (4) gives a representation of f as a convex combination of these (very special) characters.

Let \mathcal{G} denote the (infinite-dimensional) vector space of formal linear combinations (with real coefficients) of partially labeled graphs. We can turn \mathcal{G} into an algebra by using G_1G_2 introduced above as the product of two generators, and then extending this multiplication to the other elements linearly. (So \mathcal{G} is the semigroup algebra of \mathcal{PLG} .) Clearly \mathcal{G} is associative and commutative, and the empty graph is a unit element.

For every finite set $S \subseteq \mathbf{Z}_+$, the set of all formal linear combinations of S -labeled graphs forms a subalgebra $\mathcal{G}(S)$ of \mathcal{G} . The graph with $|S|$ nodes labeled by S and no edges is a unit in this algebra, which we denote by U_S .

We can extend f to a linear functional on \mathcal{G} and define an inner product

$$\langle x, y \rangle = f(xy)$$

for $x, y \in \mathcal{G}$. By our hypothesis that f is reflection positive, this inner product is positive semidefinite, i.e., $\langle x, x \rangle \geq 0$ for all $x \in \mathcal{G}$. Indeed, if $x = \sum_{G \in \mathcal{G}} x_G G \in \mathcal{G}$ (with a finite number of nonzero terms), then

$$\langle x, x \rangle = f(xx) = \sum_{G_1, G_2 \in \mathcal{G}} x_{G_1} x_{G_2} f(G_1G_2) \geq 0,$$

since the quadratic form is of the form $x^\top M(f, k)x$ for a large enough k , and $M(f, k)$ is positive semidefinite.

Let

$$\mathcal{K} = \{x \in \mathcal{G} : f(xy) = 0 \ \forall y \in \mathcal{G}\}$$

be the annihilator of \mathcal{G} . It follows from the positive semidefiniteness of the inner product that $x \in \mathcal{K}$ could also be characterized by $f(xx) = 0$.

Clearly, \mathcal{K} is an ideal in \mathcal{G} , so we can form the quotient algebra $\hat{\mathcal{G}} = \mathcal{G}/\mathcal{K}$. We can also define $\hat{\mathcal{G}}(S) = \mathcal{G}(S)/\mathcal{K}$. It is easy to check that every graph U_S has the same image under this factorization, namely the unit element u of $\hat{\mathcal{G}}$. Furthermore, if $x \in \mathcal{K}$, then $f(x) = f(xu) = 0$, and so f can also be considered as a linear functional on $\hat{\mathcal{G}}$. We denote by \hat{G} the element of $\hat{\mathcal{G}}$ corresponding to the partially labeled graph G .

Next, note that $\hat{\mathcal{G}}(S)$ is a finite-dimensional commutative algebra of dimension $r_{|S|}$, with the positive definite inner product $\langle x, y \rangle = f(xy)$.

Let $S \subseteq T$ be finite subsets of \mathbf{Z}_+ . Then $\hat{\mathcal{G}}(S) \subseteq \hat{\mathcal{G}}(T)$. Indeed, every S -labeled graph G can be turned into a T -labeled graph G' by adding $|T \setminus S|$ new isolated nodes and labeling them by the elements of $T \setminus S$. It is straightforward to check that $G - G' \in \mathcal{K}$, and so G and G' correspond to the same element of $\hat{\mathcal{G}}$.

We will also need the orthogonal projection π_S of $\hat{\mathcal{G}}$ to the subalgebra $\hat{\mathcal{G}}(S)$. This has a very simple combinatorial description. For every partially labeled graph G and $S \subseteq \mathbf{Z}_+$, let G_S denote the S -labeled graph obtained by deleting the labels not in S ; then $\pi_S(\hat{G}) = \hat{G}_S$.

To see this, let G be any partially labeled graph. Then

$$\langle G_S, G - G_S \rangle = f(G_S(G - G_S)) = f(G_S G) - f(G_S G_S) = 0$$

since GG_S and $G_S G_S$ are isomorphic graphs. Hence

$$\langle \hat{G}_S, \hat{G} - \hat{G}_S \rangle = 0,$$

showing that \hat{G}_S is indeed the orthogonal projection of \hat{G} onto $\hat{\mathcal{G}}(S)$.

We'll be interested in the idempotent elements of $\hat{\mathcal{G}}$. If p is idempotent, then

$$f(p) = f(pp) = \langle p, p \rangle > 0.$$

For two idempotents p and q , we say that q *resolves* p if $pq = q$. It is clear that this relation is transitive.

Let S be a finite subset of \mathbf{Z}_+ and set $r = r_{|S|}$. Since the algebra $\hat{\mathcal{G}}(S)$ is finite dimensional and commutative, and all its elements are self-adjoint with respect to the positive definite inner product $\langle \cdot, \cdot \rangle$, it has a (uniquely determined) basis $\mathcal{P}_S = \{p_1^S, \dots, p_r^S\}$ such that $(p_i^S)^2 = p_i^S$ and $p_i^S p_j^S = 0$ for $i \neq j$. We denote by $\mathcal{P}_{T,p}$ the set of all idempotents in \mathcal{P}_T that resolve a given idempotent p . If $|T| = |S| + 1$, then the number of elements in $\mathcal{P}_{T,p}$ will be called the *degree* of $p \in \mathcal{P}_S$ and denoted by $\deg(p)$. Obviously this value is independent of which $(|S| + 1)$ -element superset T of S we are considering.

Claim 4.1. Let r be any idempotent element of $\mathcal{G}(S)$. Then r is the sum of those idempotents in \mathcal{P}_S that resolve it.

Indeed, we can write

$$r = \sum_{p \in \mathcal{P}_S} \mu_p p$$

with some scalars μ_p . Now using that r is idempotent:

$$r = r^2 = \sum_{p, p' \in \mathcal{P}_S} \mu_p \mu_{p'} p p' = \sum_{p \in \mathcal{P}_S} \mu_p^2 p,$$

which shows that $\mu_p^2 = \mu_p$ for every p , and so $\mu_p \in \{0, 1\}$. So r is the sum of some subset $X \subseteq \mathcal{P}_S$. It is clear that $rp = p$ for $p \in X$ and $rp = 0$ for $p \in \mathcal{P}_S \setminus X$, so X consists of exactly those elements of \mathcal{P}_S that resolve q .

As a special case, we see that

$$(8) \quad u = \sum_{p \in \mathcal{P}_S} p$$

is the unit element of $\hat{\mathcal{G}}(S)$ (this is the image of the graph U_S) and also the unit element of the whole algebra $\hat{\mathcal{G}}$.

Claim 4.2. Let $S \subset T$ be two finite sets. Then every $q \in \mathcal{P}_T$ resolves exactly one element of \mathcal{P}_S .

Indeed, we have by (8) that

$$u = \sum_{p \in \mathcal{P}_S} p = \sum_{p \in \mathcal{P}_S} \sum_{\substack{q \in \mathcal{P}_T \\ q \text{ resolves } p}} q,$$

and also

$$u = \sum_{q \in \mathcal{P}_T} q,$$

so by the uniqueness of the representation we get that every q must resolve exactly one p .

Claim 4.3. Let S, T, U be finite sets and $S = T \cap U$. If $x \in \hat{\mathcal{G}}(T)$ and $y \in \hat{\mathcal{G}}(U)$, then

$$f(xy) = f(\pi_S(x)y).$$

Indeed, for every T -labeled graph G_1 and U -labeled graph G_2 , the graphs $G_1 G_2$ and $\pi_S(G_1) G_2$ are isomorphic. Hence the claim follows by linearity.

Claim 4.4. If $p \in \mathcal{P}_S$ and q resolves p , then

$$\pi_S(q) = \frac{f(q)}{f(p)} p.$$

We show that both sides give the same inner product with every basis element in \mathcal{P}_S . Since q does not resolve any $p' \in \mathcal{P}_S \setminus \{p\}$, we have $p'q = 0$ for every such p' . By Claim 4.3, this implies that

$$\langle p', \pi_S(q) \rangle = f(p' \pi_S(q)) = f(p'q) = 0 = \langle p', \frac{f(q)}{f(p)} p \rangle.$$

Furthermore,

$$\langle p, \pi_S(q) \rangle = f(p \pi_S(q)) = f(pq) = f(q) = \langle p, \frac{f(q)}{f(p)} p \rangle.$$

This proves the claim.

Claim 4.5. Let S, T, U be finite sets and $S = T \cap U$. Then for any $p \in \mathcal{P}_S$, $q \in \mathcal{P}_{T,p}$ and $r \in \hat{\mathcal{G}}(U)$ we have

$$f(p)f(qr) = f(q)f(pr).$$

Indeed, by Claims 4.4 and 4.3,

$$f(qr) = f(\pi_S(q)r) = \frac{f(q)}{f(p)} f(pr).$$

Claim 4.6. If both $q \in \mathcal{P}_T$ and $r \in \mathcal{P}_U$ resolve p , then $qr \neq 0$.

Indeed, by Claim 4.5,

$$f(qr) = \frac{f(q)}{f(p)} f(pr) = \frac{f(q)}{f(p)} f(r) > 0.$$

Claim 4.7. If $S \subset T$, and $q \in \mathcal{P}_T$ resolves $p \in \mathcal{P}_S$, then $\deg(q) \geq \deg(p)$.

It suffices to show this in the case when $|T| = |S| + 1$. Let $U \subset \mathbf{Z}_+$ be any $(|S| + 1)$ -element superset of S different from T . Let Y be the set of elements in \mathcal{P}_U resolving p . Then $p = \sum_{r \in Y} r$ by Claim 4.1. Here $|Y| = \deg(p)$. Furthermore, we have

$$\sum_{r \in Y} rq = q \sum_{r \in Y} r = qp = q.$$

Each of the terms on the left-hand side is nonzero by Claim 4.6, and since the terms are all idempotent, each of them is the sum of one or more elements of $\mathcal{P}_{T \cup U}$. Furthermore, if $r, r' \in Y$ ($r \neq r'$), then we have the orthogonality relation

$$(rq)(r'q) = (rr')q = 0,$$

so the basic idempotents in the expansion of each term are different. So the expansion of q in $\mathcal{P}_{T \cup U}$ contains at least $|Y| = \deg(p)$ terms, which we wanted to prove.

4.2. Bounding the expansion. From now on, we assume that there is a $q > 0$ such that $r_k \leq q^k$ for all k .

So if a basic idempotent $p \in \mathcal{P}_S$ has degree D , then there are D basic idempotents on the next level with degree $\geq D$, and hence if $|T| \geq |S|$, then the dimension of $\hat{\mathcal{G}}(T)$ is at least $D^{|T \setminus S|}$. It follows that the degrees of basic idempotents are bounded by q ; let D denote the maximum degree, attained by some $p \in \mathcal{P}_S$.

Let us fix such a set S and $p \in \mathcal{P}_S$ with maximum degree D . For $u \in \mathbf{Z}_+ \setminus S$, let q_1^u, \dots, q_D^u denote the elements of $\mathcal{P}_{S \cup \{u\}}$ resolving p . Note that for $u, v \in \mathbf{Z}_+ \setminus S$, there is a natural isomorphism between $\hat{\mathcal{G}}(S \cup \{u\})$ and $\hat{\mathcal{G}}(S \cup \{v\})$ (induced by the map that fixes S and maps u onto v), and we may choose the labeling so that q_i^u corresponds to q_i^v under this isomorphism.

Next we describe, for a finite set $T \supset S$, all basic idempotents in \mathcal{P}_T that resolve p . Let $V = T \setminus S$, and for every map $\phi : V \rightarrow \{1, \dots, D\}$, let

$$(9) \quad q_\phi = \prod_{v \in V} q_{\phi(v)}^v.$$

Note that by Claim 4.5,

$$(10) \quad f(q_\phi) = f\left(\prod_{v \in V} q_{\phi(v)}^v\right) = \left(\prod_{v \in V} \frac{f(q_{\phi(v)}^v)}{f(p)}\right) f(p) \neq 0,$$

and so $q_\phi \neq 0$.

Claim 4.8.

$$\mathcal{P}_{T,p} = \{q_\phi : \phi \in \{1, \dots, D\}^V\}.$$

We prove this by induction on $|T \setminus S|$. For $|T \setminus S| = 1$ the assertion is trivial. Suppose that $|T \setminus S| > 1$. Let $u \in T \setminus S$, $U = S \cup \{u\}$ and $W = T \setminus \{u\}$; thus $U \cap W = S$. By the induction hypothesis, the basic idempotents in \mathcal{P}_W resolving p are elements of the form q_ψ ($\psi \in \{1, \dots, D\}^{V \setminus \{u\}}$).

Let r be one of these. By Claim 4.6, $rq_i^u \neq 0$ for any $1 \leq i \leq D$ and clearly resolves r . We can write rq_i^u as the sum of basic idempotents in \mathcal{P}_T , and it is easy to see that these also resolve r . Furthermore, the basic idempotents occurring in

the expression of rq_i^u and rq_j^u ($i \neq j$) are different. But r has degree D , so each rq_i^u must be a basic idempotent in \mathcal{P}_T itself.

Since the sum of the basic idempotents rq_i^u ($r \in \mathcal{P}_{W,p}$, $1 \leq i \leq D$) is p , it follows that these are all the elements of $\mathcal{P}_{T,p}$. This proves the claim.

It is immediate from the definition that an idempotent q_ϕ resolves q_i^v if and only if $\phi(v) = i$. Hence it also follows that

$$(11) \quad q_i^v = \sum_{\phi: \phi(v)=i} q_\phi.$$

4.3. Constructing the target graph. Now we can define H as follows. Let H be the looped complete graph on $V(H) = \{1, \dots, D\}$. We have to define the node weights and edge weights.

Fix any $u \in \mathbf{Z}_+ \setminus S$. For every $i \in V(H)$, let

$$\alpha_i = \frac{f(q_i^u)}{f(p)}$$

be the weight of the node i . Clearly $\alpha_i > 0$.

Let $u, v \in \mathbf{Z}_+ \setminus S$, $v \neq u$, and let $W = S \cup \{u, v\}$. Let K_{uv} denote the graph on W which has only one edge connecting u and v , and let k_{uv} denote the corresponding element of $\hat{\mathcal{G}}(W)$. We can express pk_{uv} as a linear combination of elements of $\mathcal{P}_{W,p}$ (since for any $r \in \mathcal{P}_W \setminus \mathcal{P}_{W,p}$ one has $rp = 0$ and hence $rp k_{u,v} = 0$):

$$pk_{uv} = \sum_{i,j} \beta_{ij} q_i^u q_j^v.$$

This defines the weight β_{ij} of the edge ij . Note that $\beta_{ij} = \beta_{ji}$ for all i, j , since $pk_{uv} = pk_{vu}$.

We prove that this weighted graph H gives the right homomorphism function.

Claim 4.9. For every finite graph G , $f(G) = \text{hom}(G, H)$.

By (11), we have for each pair u, v of distinct elements of $V(G)$,

$$pk_{uv} = \sum_{i,j \in V(H)} \beta_{i,j} q_i^u q_j^v = \sum_{i,j \in V(H)} \beta_{i,j} \sum_{\phi: \phi(u)=i, \phi(v)=j} q_\phi = \sum_{\phi \in V(H)^V} \beta_{\phi(u), \phi(v)} q_\phi.$$

Consider now any V -labeled graph G with $V(G) = V$, and let g be the corresponding element of $\hat{\mathcal{G}}$. Then

$$\begin{aligned} pg &= \prod_{uv \in E(G)} pk_{uv} = \prod_{uv \in E(G)} \left(\sum_{\phi \in V(H)^V} \beta_{\phi(u), \phi(v)} q_\phi \right) \\ &= \sum_{\phi: V \rightarrow V(H)} \left(\prod_{uv \in E(G)} \beta_{\phi(u), \phi(v)} \right) q_\phi. \end{aligned}$$

Since $p \in \hat{\mathcal{G}}(S)$, $g \in \hat{\mathcal{G}}(V)$ and $S \cap V = \emptyset$, we have $f(p)f(g) = f(pg)$, and so by (10),

$$\begin{aligned} f(p)f(g) &= f(pg) = \sum_{\phi \in V(H)^V} \left(\prod_{uv \in E(G)} \beta_{\phi(u), \phi(v)} \right) f(q_\phi) \\ &= \sum_{\phi: V \rightarrow V(H)} \left(\prod_{uv \in E(G)} \beta_{\phi(u), \phi(v)} \right) \left(\prod_{v \in V(G)} \alpha_{\phi(v)} \right) f(p). \end{aligned}$$

The factor $f(p) > 0$ can be cancelled from both sides, completing the proof.

5. EXTENSIONS: GRAPHS WITH LOOPS, DIRECTED GRAPHS AND HYPERGRAPHS

In our arguments we allowed parallel edges in G , but no loops. Indeed, the representation theorem is false if G can have loops: it is not hard to check that the graph parameter

$$f(G) = 2^{\#\text{loops}}$$

cannot be represented as a homomorphism function, even though its connection matrix $M(f, k)$ is positive semidefinite and has rank 1. To get a representation theorem for graphs with loops, each loop e in the target graph H must have two weights: one which is used when a nonloop edge of G is mapped onto e , and the other, when a loop of G is mapped onto e . With this modification, the proof goes through.

The constructions and results above are in fact more general; they extend to directed graphs and hypergraphs. There is a common generalization to these results, using semigroups. This will be stated and proved in a separate paper. An analogous result for edge coloring models (with a substantially more difficult proof) was proved by B. Szegedy [8].

ACKNOWLEDGEMENT

We are indebted to Christian Borgs, Jennifer Chayes, Monique Laurent, Miki Simonovits, Vera T. Sós, Balázs Szegedy, Gábor Tardos and Kati Vesztegombi for many valuable discussions and suggestions on the topic of graph homomorphisms.

REFERENCES

- [1] C. Berg, J.P.R. Christensen, P. Ressel: Positive definite functions on abelian semigroups, *Mathematische Annalen* **223** (1976), 253–272. MR0420150 (54:8165)
- [2] C. Berg, J.P.R. Christensen, P. Ressel: *Harmonic Analysis on Semigroups—Theory of Positive Definite and Related Functions*, Graduate Texts in Mathematics **100**, Springer, New York, 1984. MR0747302 (86b:43001)
- [3] C. Berg and P.H. Maserick: Exponentially bounded positive definite functions, *Illinois Journal of Mathematics* **28** (1984), 162–179. MR0730718 (85i:43012)
- [4] M. Freedman, L. Lovász and D. Welsh (unpublished).
- [5] P. de la Harpe and V.F.R. Jones: Graph Invariants Related to Statistical Mechanical Models: Examples and Problems, *Journal of Combinatorial Theory B* **57** (1993), 207–227. MR1207488 (94c:05033)
- [6] R.J. Lindahl and P.H. Maserick: Positive-definite functions on involution semigroups, *Duke Mathematical Journal* **38** (1971), 771–782. MR0291826 (45:916)
- [7] L. Lovász, B. Szegedy: Limits of graph sequences, Microsoft Research Technical Report MSR-TR-2004-79, <ftp://ftp.research.microsoft.com/pub/tr/TR-2004-79.pdf>
- [8] B. Szegedy: Edge coloring models and reflection positivity (manuscript), <http://arxiv.org/abs/math.CO/0505035>

- [9] D.J.A. Welsh: *Complexity: Knots, Colourings and Counting*, London Math. Soc. Lecture Note Series **186**, Cambridge University Press, Cambridge, 1993. MR1245272 (94m:57027)
- [10] H. Whitney: The coloring of graphs, *Ann. of Math.* **33** (1932), 688-718. MR1503085

MICROSOFT INSTITUTE FOR QUANTUM PHYSICS, SANTA BARBARA, CALIFORNIA 93106

MICROSOFT RESEARCH, ONE MICROSOFT WAY, REDMOND, WASHINGTON 98052

CWI, KRUISLAAN 413, 1098 SJ AMSTERDAM, THE NETHERLANDS