Contractors and connectors of graph algebras[*]

LÁSZLÓ LOVÁSZ and BALÁZS SZEGEDY
Microsoft Research
One Microsoft Way

April 2005

Microsoft Research Technical Report # MSR-TR-2005-91

Abstract

We study generalizations of the "contraction-deletion" relation of the Tutte polynomial, and other similar simple operations, to other graph parameters. The question can be set in the framework of graph algebras introduced by Freedman, Lovász and Schrijver in [2], and it relates to their behavior under basic graph operations like contraction and subdivision.

Graph algebras were introduced in [2] to study and characterize homomorphism functions. We prove that for homomorphism functions, these graph algebras have special elements called "contractors" and "connectors". This gives a new characterization of homomorphism functions.

1   Introduction and results

The contraction-deletion operation for the Tutte polynomial is a basic tool in graph theory. For our purposes, let us formulate this as property as follows: Let G be a graph and u and v nonadjacent nodes in G. Let G' be obtained by identifying these nodes. Then the Tutte polynomial of G' can be expressed as a linear combination of the Tutte polynomials of G and the graph $G + uv$ (obtained by connecting u and v by an edge).

Which other graph parameters have a similar property that the parameter of G' can be expressed as a linear combination of the parameter on graphs obtained from G by attaching various "small" graphs at u and v?

---

[*]AMS Subject Classification: Primary 05C99, Secondary 16S99

If we study the number of perfect matchings in a graph, then a useful observation is that subdividing an edge by two new nodes does not change this number. Which other graph parameters have a similar property that the parameter of G can be expressed as a linear combination of the parameter on graphs obtained from G by deleting the edge uv and attaching various "small" graphs at u and v?

These questions are related to the work in [2] and subsequent work [3, 5]. Here certain algebras generated by graphs played a useful role, and the above questions can be stated as rather basic properties of these algebras. Among others, they can phrased in terms of the existence of special elements called "contractors" and "connectors".

Graph algebras were introduced in [2] to study and characterize homomorphism functions. We prove that for homomorphism functions, these graph algebras have contractors and connectors. This gives a new characterization of homomorphism functions.

1.1 Graph algebras

To state our results, we need to introduce some formalism. Fix a positive integer k. A *k-labeled graph* is a finite graph in which some of the nodes are labeled by numbers $1, \ldots, k$. We denote by K_k the k-labeled complete graph with k nodes, and by O_k, the k-labeled graph with k nodes and no edges. A k-labeled quantum graph is *simple*, if it has no multiple edges, and its labeled nodes are independent.

Let F_1 and F_2 be two k-labeled graphs. Their *product* F_1F_2 is defined as follows: we take their disjoint union, and then identify nodes with the same label. For two $\emptyset$-labeled graphs, F_1F_2 is their disjoint union. Clearly this multiplication is associative and commutative.

A *k-labeled quantum graph* is a formal finite linear combination (with real coefficients) of k-labeled graphs. Let $\mathcal{G}_k$ denote the (infinite dimensional) vector space of all k-labeled quantum graphs. We can turn $\mathcal{G}_k$ into an algebra by using F_1F_2 introduced above as the product of two generators, and then extending this multiplication to the other elements linearly. Clearly $\mathcal{G}_k$ is associative and commutative. The graph O_k is a unit element in $\mathcal{G}_k$.

We'll also consider the subalgebra $\mathcal{G}_k^{\mathrm{simp}}$ generated by k-labeled simple graphs and the subalgebra $\mathcal{G}_k^0$ generated by those k-labeled graphs whose labeled points are independent. It is clear that $\mathcal{G}_k^{\mathrm{simp}} \subseteq \mathcal{G}_k^0 \subseteq \mathcal{G}_k$.

1.2   Graph parameters and star algebras

A *graph parameter* is a function defined on graphs, invariant under isomorphism. We call a graph parameter f *multiplicative* if for any two (0-labeled) graphs F_1, F_2 we have

$$f(F_1 F_2) = f(F_1) f(F_2).$$

Every graph parameter f introduces further structure on the algebras $\mathcal{G}_k$. We extend f linearly to quantum graphs, and consider $f(x)$ as a "trace" of x. We use this trace function to introduce an inner product on $\mathcal{G}$ by

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

Let $\mathcal{N}_k(f)$ denote the kernel of this inner product, i.e.,

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

Then we can define the factor algebra

$$\mathcal{G}_k/f = \mathcal{G}_k/\mathcal{N}_k(f).$$

For $x, y \in \mathcal{G}_k$, we write $x \equiv y \pmod{f}$ if $x - y \in \mathcal{N}_k(f)$.

The dimension of $\mathcal{G}_k/f$ is called the *rank-connectivity* of the parameter f, and is denoted by $\mathrm{rk}(f, k)$. This is in general infinite, but it is finite for quite a few interesting graph parameters.

We say that f is *reflection positive* if this inner product is semidefinite: $\langle x, x \rangle \geq 0$ for all x. In this case, we have

$$\mathcal{N}_k(f) = \{x \in \mathcal{G}_k : \ f(x^2) = 0\}.$$

Furthermore the algebra $\mathcal{G}_k/f$ is a commutative algebra whose elements form a Hilbert space with the property that $\langle xy, z \rangle = \langle x, yz \rangle$. If this Hilbert space is finite dimensional, then this implies that $\mathcal{G}_k/f$ has a basis $p_1, \ldots, p_N$ such that $p_i^2 = p_i$ and $p_i p_j = 0$ for $i \neq j$.

In this paper we prove a number of facts relating these graph algebras to the basic graph operations of subdivision and contraction. This will lead to alternate characterizations of "homomorphism functions" defined in Section 1.3.

3

1.3 Homomorphism functions

This important class of graph parameters was the motivating example for the studies in this paper. For two graphs F and G, let $\hom(F, G)$ denote the number of homomorphisms (adjacency preserving maps) from $V(F)$ to $V(G)$.

We need to generalize this to the case when G is weighted. A *weighted graph* G is a graph with a weight $\alpha_G(i)$ associated with each node and a weight $\beta_G(i, j)$ associated with each edge ij. In this paper we assume that the nodeweights are positive. Let $\alpha_G = \sum_{i \in V(G)}$ denote the total nodeweight of G. (An unweighted graph can be considered as a weighted graph where all the node- and edgeweights are 1.)

To every function $\phi : S \to V(G)$ with $S \subseteq V(F)$ we assign the weight

$$\hom_\phi(F, G) = \sum_{\psi: V(F) \to V(G),\ \psi|_S = \phi} \prod_{i \in V(F) \setminus S} \alpha_G(\phi(i)) \prod_{u,v \in V(F)} \beta_G(\psi(u), \psi(v)).$$

We then define

$$\hom(F, G) = \hom_\emptyset(F, G).$$

Sometimes it is more convenient to use the "homomorphism density"

$$t(F, G) = \frac{\hom(F, G)}{\alpha_G^{|V(F)|}}.$$

If G is unweighted, this specializes to

$$t(F, G) = \frac{\hom(F, G)}{|V(G)|^{|V(F)|}}.$$

We'll also consider the number $\mathrm{inj}(F, G)$ of injective homomorphisms of F into G, and its normalized version

$$t_0(F, G) = \frac{\mathrm{inj}(F, G)}{(|V(G)|)_{|V(F)|}}$$

(where $(n)_k = n(n-1) \ldots (n-k+1)$).

The following theorem was proved in [2]:

Theorem 1.1 *A graph parameter f can be represented in the form $f = \hom(., H)$ for some finite weighted graph H on at most q nodes if and only if it is reflection positive and $\mathrm{rk}(f, k) \leq q^k$ for all $k \geq 0$.*

An exact formula for $\mathrm{rk}(f, k)$ for homomorphism functions was obtained in [3]. To state it, we need a definition. Two nodes i and j in a weighted graph H are *twins*, if $\beta_H(i, k) = \beta_H(j, k)$ for every node $k \in V(H)$. Twin nodes can be merged without changing the homomorphism functions $t(., H)$ and $\hom(., H)$.

Theorem 1.2 *If $f = \hom(., H)$, and H has no twin nodes, then $\mathrm{rk}(f, k)$ is the number of orbits of the automorphism group of H on the ordered k-tuples of nodes.*

1.4   Contractors and connectors

For a 2-labeled graph F in which the two labeled nodes are nonadjacent, let F' denote the graph obtained by identifying the two labeled nodes. The map $F \mapsto F'$ maps 2-labeled graphs to 1-labeled graphs. We can extend it linearly to get an algebra homomorphism $x \mapsto x'$ from $\mathcal{G}_2^0$ into $\mathcal{G}_1$.

The map $x \mapsto x'$ does not in general preserve the inner product or even its kernel; we say that the graph parameter f is *contractible*, if for every $x \in \mathcal{G}_2^0$, $x \equiv 0 \pmod{f}$ implies $x' \equiv 0 \pmod{f}$; in other words, $x \mapsto x'$ factors to a linear map $\mathcal{G}_2^0/f \to \mathcal{G}_1/f$.

We say that $z \in \mathcal{G}_2$ is a *contractor for* f if for every $x \in \mathcal{G}_2^0$, we have

$$f(xz) = f(x').$$

Informally, attaching z at two nodes acts like identifying those two nodes.

Our second concern is to get rid of multiple edges. We say that $z \in \mathcal{G}_2^{\mathrm{simp}}$ is a *connector* for f, if $z \equiv K_2 \pmod{f}$, i.e., for every $x \in \mathcal{G}_2$ we have

$$f(zx) = f(K_2 x).$$

Note that K_2 is always a connector, but it is not simple in the sense defined above.

If a graph parameter f has a simple connector z, then every k-labeled quantum graph is congruent to a simple quantum graph modulo f. Indeed, for every $x \in \mathcal{G}_k$, in every k-labeled graph in the expansion of x, every edge can be replaced by the simple connector, which creates a simple quantum graph. In other words, $\mathcal{G}_k^{\mathrm{simp}}/f = \mathcal{G}_k/f$.

Several general facts about connectors and contractors will be stated and proved in Section 3.

1.5 The algebra of concatenations

For two 2-labeled graphs F_1 and F_2, we define their *concatenation* by identifying node 2 of F_1 with node 1 of F_2, and unlabeling this new node. We denote the resulting 2-labeled graph by $F_1 \circ F_2$. It is easy to check that this operation is associative (but not commutative). We extend this operation linearly over $\mathcal{G}_2$.

This algebra has a $*$ operation: for a 2-labeled graph F, we define F^* by interchanging the two labels. Clearly $(F_1 F_2)^* = F_2^* F_1^*$. We can also extend this linearly over $\mathcal{G}_2$.

Let f be a graph parameter. It is easy to see that if $x \equiv 0 \pmod{f}$ then $x^* \equiv 0 \pmod{f}$, so the $*$ operator is well defined on elements of $\mathcal{G}_2/f$. An further important property of concatenation is that

$$f((x \circ y)z) = f(x(z \circ y^*))$$

for any three elements $x, y, z \in \mathcal{G}_2$. It follows that if $x \equiv 0 \pmod{f}$ then $x \circ y \equiv 0 \pmod{f}$ for every $y \in \mathcal{G}_2$ and thus concatenation is also well defined on the elements of $\mathcal{G}_2/f$. It is easy to see that $\mathcal{A}_1 = (\mathcal{G}_2/f, +, \circ)$ is an associative (but not necessarily commutative) algebra. Note that if $x, y \in \mathcal{G}_2$ then $x \circ y \in \mathcal{G}_2^0$. It follows in particular that if $\mathcal{A}_1$ has a unit element then $\mathcal{G}_2^0/f = \mathcal{G}_2/f$.

Lemma 1.3 *If an element $z \in \mathcal{G}_2$ is a contractor for f then the image of z under the map $\mathcal{G}_2 \to \mathcal{G}_2/f$ is the unit element of the algebra $\mathcal{A}_1$.*

Proof. We have to check that $z \circ x \equiv x \pmod{f}$ for all $x \in \mathcal{G}_2$. This is equivalent with $f((z \circ x)y) = f(xy)$ for all $x, y \in \mathcal{G}_2$. Using that $x \circ y \in \mathcal{G}_2^0$ we obtain that

$$f((z \circ x)y) = f(z(y \circ x^*)) = f((y \circ x^*)') = f(xy).$$

$\square$

1.6 Contractors and connectors for homomorphism functions

The first two results in this paper concerns graph parameters that are homomorphism functions, i.e., they are of the form $f = \hom(., H)$ for some weighted graph H. It is easy to check from the definitions that a 2-labeled quantum graph z is a contractor for $\hom(., G)$ if and only if

$$\hom_\phi(z, G) = \begin{cases} 1 & \text{if } \phi(1) = \phi(2), \\ 0 & \text{otherwise,} \end{cases} \tag{1}$$

for every $\phi : \{1, 2\} \to V(G)$. It is a connector for $\mathrm{hom}(., G)$ if and only if $z \in \mathcal{G}_2^{\mathrm{simp}}$, and

$$\mathrm{hom}_\phi(z, G) = \beta_G(\phi(1), \phi(2)) \tag{2}$$

for every $\phi : \{1, 2\} \to V(G)$.

We denote by P_n the path with n nodes, with the two endnodes labeled 1 and 2 (so $P_2 = K_2$). A *quantum path* is a linear combination of such paths. A *series-parallel graph* is a 2-labeled graph obtained from K_2 by repeated application of the product and concatenation operations. A *series-parallel quantum graph* is a linear combination of series-parallel graphs.

Theorem 1.4 *Let $f = \mathrm{hom}(., H)$ for some finite weighted graph H. Then f has a contractor and also a simple connector. Furthermore, it has a contractor that is a series-parallel quantum graph and a simple connector that is a quantum path.*

Using the notion of a contractor, we can give the following characterization of homomorphism functions (we don't know whether a similar theorem holds using some special connectors instead of contractors).

Theorem 1.5 *A graph parameter f can be represented in the form $f = \mathrm{hom}(., H)$ for some finite weighted graph H if and only if it is multiplicative, reflection positive and has a contractor.*

1.7 Homomorphisms into measure graphs

Every symmetric measurable function $W : [0, 1]^2 \to [0, 1]$ defines a graph parameter $t(., W)$ by

$$t(F, W) = \int\limits_{[0,1]^n} \prod_{ij \in E(F)} W(x_i, x_j) \, dx_1 \ldots dx_n.$$

Homomorphism density functions into finite weighted graphs is a special case. Call a symmetric function $W : [0, 1]^2 \to [0, 1]$ a *step function*, if there is a partition $[0, 1] = A_1 \cup \cdots \cup A_q$ into measurable sets such that W is constant on $A_i \times A_j$ for all $1 \le i, j \le q$. It is trivial to check that if W is a step function, then $t(., W) = t(., H)$ for a finite weighted graph H and vice versa. It was noted in [2] that the graph parameter $t(., W)$ is reflection positive, and it is obvious that it is multiplicative.

These parameters occur in the context of limits of graph sequences. We say that a sequence (G_n) of simple graphs is *convergent* if $t(F, G_n)$ converges

7

to some value $t(F)$ for every simple graph F. In [4] it was shown that (at least for the case of parameters defined on simple graphs), the parameters $t(.,W)$ are precisely the limits of parameters $t(F)$ obtained this way.

Theorem 1.6 *The graph parameter $t(.,W)$ is contractible, but has no contractor unless W is a step function.*

2 Examples

The following examples are described in more detail in [2]. Here we only discuss those properties of them that relate to contractors and connectors.

2.1 Matchings

Let $\mathsf{perf}(G)$ denote the number of perfect matchings in the graph G. It is trivial that $\mathsf{perf}(.)$ is multiplicative. Its node-rank-connectivity is exponentially bounded,

$$\mathrm{rk}(\mathsf{perf}, k) = 2^k,$$

but perf is not reflection-positive. Thus $\mathsf{perf}(G)$ cannot be represented as a homomorphism function.

On the other hand: perf has a contractor: a path of length 2, and also a simple connector: a path P_4 of length 3.

2.2 Chromatic polynomial

Let $\mathsf{chr}(G) = \mathsf{chr}(G, x)$ denote the chromatic polynomial of the graph G. For every fixed x, this is a multiplicative graph parameter. 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, we have [6]

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

Note that this is always finite, but if $x \notin \mathbf{Z}_+$, then it grows faster than c^k for every c. Furthermore, $M(\mathsf{chr}, k)$ is positive semidefinite if and only if either x is a positive integer or $k \leq x + 1$. The parameter $M(\mathsf{chr}, k)$ 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 $\mathsf{chr}(G, x) = \hom(G, K_x)$.

This parameter has a contractor for every x: the 2-labeled quantum graph $K_2 - O_2$ (which amounts to the standard contraction-deletion identity

for the chromatic polynomial). It is not hard to check that $\frac{1}{x-1}P_3 - \frac{x-2}{x-1}P_2$ is a simple connector if $x \neq 1$; for $x = 1$, the chromatic polynomial is 0 if there is an edge, so P_2 is a simple connector.

2.3 Flows

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 $\mathsf{flo}(G)$ be the number of S-flows. This number is independent of the orientation. In the case when $S = \Gamma \setminus \{0\}$, $\mathsf{flo}(G)$ is the number of nowhere-0 Γ-flows.

The parameter $\mathsf{flo}(G)$ can be described as a homomorphism function [2]. It has a trivial simple connector, a path of length 2 (which is an algebraic way of saying that if we subdivide an edge, then the flows don't change essentially). In the case of nowhere-0 flows, $K_2 + O_2$ is a contractor (which amounts to the contraction-deletion identity for the flow polynomial), but in general, there does not seem to be a simple explicit construction for a contractor.

2.4 Tutte polynomial

Consider the following version of the Tutte polynomial: in terms of the variables q and v, we have

$$\mathsf{tut}(G; q, v) = \sum_{A \subseteq E(G)} q^{c(A)} v^{|A|},$$

where $c(A)$ denotes the number of components of the graph $(V(G), A)$. This differs from the usual Tutte polynomial $T(x, y)$ on two counts: first, instead of the standard variables x and y, we use $q = (x-1)(y-1)$ and $v = y-1$; second, we scale by $q^{c(E)} v^{n-c(E)}$. This way we lose the covariance under matroid duality; but we gain that the contraction/deletion relation holds for all edges e:
$$\mathsf{tut}(G) = v\mathsf{tut}(G/e) + \mathsf{tut}(G \setminus e). \tag{3}$$

If i is an isolated node of G, then we have

$$\mathsf{tut}(G - i) = q\mathsf{tut}(G). \tag{4}$$

If G is the empty graph (no nodes, no edges), then $\mathsf{tut}(G) = 1$. Another way of expressing (3) is that $K_2 - vO_2$ is a contractor of tut. It is not hard to check that $(1/v)P_3 - (1 + q/v)O_2$ is a simple connector.

The chromatic polynomial and the number of nowhere-0 Γ-flows are special substitutions into the Tutte polynomial. More precisely,

$$\mathsf{chr}(G; x) = \mathsf{tut}(G; x, -1),$$

and the number of nowhere-0 k-flows is

$$\mathsf{flo}(G) = \frac{(-1)^{|E(G)|}}{k^{|V(G)|}} \mathsf{tut}(G; k, -k).$$

It can be shown [6] that for $v \neq 0$, the Tutte polynomial behaves exactly as the corresponding chromatic polynomial:

$$\mathrm{rk}(\mathsf{chr}, k) = \begin{cases} B_{kq} & \text{if } q \text{ is a nonnegative integer,} \\ B_k & \text{otherwise.} \end{cases}$$

Furthermore, $\mathsf{tut}(G; q, v)$ is reflection positive if and only if q is a positive integer. Theorem 1.1 implies that in this case $\mathsf{tut}(G; q, v)$ is a homomorphism function, while for other substitutions it is not.

2.5 The role of multiple edges

Let, for each (multi)graph G, $\tilde{G}$ denote the (simple) graph obtained from G by keeping only one copy of each parallel class of edges. Consider a random graph H on N nodes with edge probability $1/2$, then

$$\mathsf{expt}(G) = \mathsf{E}(t_0(G, H)) = 2^{-|E(\tilde{G})|}$$

is independent of N. It is not hard to see that we also have with probability 1

$$\mathsf{expt}(G) = \lim_{N \to \infty} t(G, H).$$

From this (or from direct computation) it follows that this graph parameter is multiplicative and reflection positive. It can be checked that $\mathrm{rk}(\mathsf{expt}, k) = 2^{\binom{k}{2}}$, which is finite for every k, but has superexponential growth.

The graph parameter expt is not contractible. Consider the 3-star S_4 with 2 endnodes labeled and the path P_4 with 3 edges with both endnodes labeled. Then $S_4 \equiv P_4 \pmod{f}$, but identifying the labeled nodes produces a pair of parallel edges in S_4 but not in P_4, so $f(S'_4) = 1/4$ but $f(P'_4) = 1/8$, showing that $S'_4 \not\equiv P'_4 \pmod{f}$. This implies by lemma 3.1 below that expt does not have a contractor. It is easy to see that f does not have a simple connector either.

2.6 The number of eulerian orientations

Let $\mathsf{eul}(G)$ denote the number of eulerian orientations of the graph G. It was remarked in [4] that this parameter can be expressed as $\mathsf{eul}(G) = t(G, W)$, where

$$W(x, y) = 2 \cos(2\pi(x - y)).$$

Thus it follows by Theorem 1.6 that eul is contractible, but has no contractor. It is easy to see that a path of length 2 is a simple connector.

3 General facts about connectors and contractors

We start with an easy observation.

Proposition 3.1 *If a graph parameter has a contractor, then it is contractible.*

Proof. Let z be a contractor for f. Suppose that $x \in \mathcal{G}_2$ satisfies $x \equiv 0$ (mod f), and let $y \in \mathcal{G}_1$. Choose a $\hat{y} \in \mathcal{G}_2$ such that $\hat{y}' = y$. Then

$$f(x'y) = f(x'\hat{y}') = f((x\hat{y})') = f((x\hat{y})z) = f(x(\hat{y}z)) = 0,$$

showing that $x' \equiv 0$ (mod f). $\qquad\qquad\qquad\qquad\qquad \square$

While the existence of a contractor does not imply the existence of a simple connector or vice versa, there is some connection, as expressed in the following proposition (see also Corollary 3.6).

Proposition 3.2 *If f is contractible, has a simple connector, and $\mathrm{rk}(f, 2)$ is finite, then f has a contractor.*

Proof. Since $\langle x, y \rangle = f(xy)$ is a symmetric (possibly indefinite) bilinear form that is not singular on $\mathcal{G}_2/f$, there is a basis $p_1, \dots, p_r$ in $\mathcal{G}_2/f$ such that $f(p_i p_j) = 0$ if $i \neq j$ and $f(p_i p_i) \neq 0$. By the assumption that f is simplifiable, we may represent this basis by simple quantum graphs; then the contracted quantum graphs p_i' are defined. Let

$$z = \sum_{i=1}^{N} \frac{f(p_i')}{f(p_i^2)} p_i.$$

We claim that z is a contractor. Indeed, let $x \in \mathcal{G}_2^{\mathrm{simp}}$, and write

$$x \equiv \sum_{i=1}^{N} a_i p_i \quad (\text{mod } f).$$

Then we have

$$f(xz) = \sum_{i=1}^{N} a_i \frac{f(p_i')}{f(p_i^2)} f(p_i^2) = \sum_{i=1}^{N} f(p_i')a_i.$$

On the other hand, contractibility implies that

$$x' \equiv \sum_{i=1}^{N} a_i p_i' \pmod{f},$$

and so

$$f(x') = \sum_{i=1}^{N} a_i f(p_i') = f(xz).$$

$\square$

Proposition 3.3 *If $M(f,2)$ is positive semidefinite, f is contractible and $\mathrm{rk}(f,2)$ is finite, then f has a simple connector that is a quantum path.*

For the proof, we need the following simple lemma.

Lemma 3.4 *Assume that $M(f,2)$ is positive semidefinite. Let $x \in \mathcal{G}_2$, and assume that $x \circ P_3 \equiv 0 \pmod{f}$. Then $x \circ P_2 \equiv 0 \pmod{f}$.*

Proof. We have
$$(x \circ P_2)^2 = (x \circ P_3)x = 0,$$
and by reflection positivity, this implies that $x \circ P_2 \equiv 0 \pmod{f}$. $\square$

Proof [of Proposition 3.3]: Since $\mathcal{G}_2/f$ is finite dimensional, there is a linear dependence between $P_2, P_3, \ldots$ in $\mathcal{G}_2/f$. Hence there is a (smallest) $k \geq 2$ such that P_k can be expressed as

$$P_k \equiv \sum_{i=1}^{N} a_i P_{k+i} \pmod{f} \tag{5}$$

with some positive integer N and real numbers $a_1, \ldots, a_N$. The assertion is equivalent to saying that $k = 2$.

Let $x = P_2 - \sum_{i=1}^{N} a_i P_{2+i}$. Then (5) can be written as $x \circ P_{k-1} \equiv 0 \pmod{f}$. If $k > 3$, then Lemma 3.4 implies that $x \circ P_{k-2} \equiv 0 \pmod{f}$, which contradicts the minimality of k. Suppose that $k = 3$. Then from (5) we have that $(x - \sum_{i=1}^{N} a_i x \circ P_{1+i}) \circ P_3 \equiv x \circ x \circ P_2 \equiv 0 \pmod{f}$. By Lemma 3.4 we get that $x \circ x \equiv 0 \pmod{f}$ and using contractibility we obtain that $0 = f((x \circ x)') = f(x^2)$. Now reflection positivity shows that $x \equiv 0 \pmod{f}$.

Corollary 3.5 *If $M(f,2)$ is positive semidefinite, f is contractible and* $\mathrm{rk}(f,2)$ *is finite, then* $\mathcal{G}_k = \mathcal{G}_k^{\mathrm{simp}}$ *for every $k \geq 1$.*

The following statement is a corollary of Proposition 3.3 and Proposition 3.2.

Corollary 3.6 *If $M(f,2)$ is positive semidefinite, f is contractible and* $\mathrm{rk}(f,2)$ *is finite then f has a contractor.*

4 Homomorphism functions: proofs

4.1 Proof of Theorem 1.4

Suppose that $f = \hom(.,G)$ for some weighted graph G. We may assume that G is twin-free.

We start with constructing a connector. Let $\alpha_1, \ldots, \alpha_m$ be the nodeweights and β_{ij} $(i,j = 1, \ldots, m)$, the edgeweights of G. Let $B = (\beta_{ij})$ be the (weighted) adjacency matrix of G, and let $D = \mathrm{diag}(\sqrt{\alpha_1}, \ldots, \sqrt{\alpha_m})$. Let $\lambda_1, \ldots, \lambda_t$ be the nonzero eigenvalues of the matrix DBD (which are real as DBD is symmetric), and consider the polynomial $\rho(z) = z \prod_{i=1}^{t}(1 - z/\lambda_i)$. Then $\rho(DBD) = 0$. Since the constant term in $\rho(z)$ is 0 and the linear term is z, this expresses DBD as a linear combination of higher powers of DBD:

$$DBD = \sum_{s=2}^{t} a_s (DBD)^s,$$

or

$$B = \sum_{s=2}^{t} a_s (BD^2)^{s-1} B. \tag{6}$$

For every mapping $\varphi : \{1,2\} \to V(G)$, we have

$$\hom_\varphi(P_s, G) = ((BD^2)^{s-2} B)_{\varphi(1)\varphi(2)}.$$

Let

$$y = \sum_{s=2}^{t} a_s P_{s+1},$$

Then (6) implies that for every 2-labeled graph G,

$$f(K_2 G) = f(yG).$$

Thus y is a connector. By construction, it is a linear combination of paths.

13

For the existence of a contractor, there are two general arguments.

First, we can use Lemma 3.2: it is easy to check that f is contractible; the condition that $\mathcal{G}_1/f = \mathcal{G}_2'/f$ follows from the existence of a connector; and $M(f,2)$ is positive semidefinite and has finite finite rank by Theorem 1.1.

Second, to prove that there exists a 2-labeled quantum graph z satisfying (1), we can invoke the following result [3]:

Theorem 4.1 *Let G be a twin-free weighted graph and $\Phi : V(G)^k \to \mathbf{R}$. Then there exists a k-labeled quantum graph z such that*

$$\hom_\phi(z, G) = \Phi(\phi)$$

for every $\phi \in V(G)^k$, if and only if Φ is invariant under the automorphisms of G: for every $\phi \in V(H)^k$ and every automorphism σ of H, $\Phi(\phi\sigma) = \Phi(\phi)$.

However, it is worth while to give a third, more specific argument, because it gives the stronger result that a series-parallel contractor exists. To every 2-labeled quantum graph x, we assign the $V(G) \times V(G)$ matrix $M(x)$ as follows: for $i, j \in V(G)$, we set

$$M(x)_{ij} = \hom_{1 \mapsto i, 2 \mapsto j}(x, G).$$

Then it is easy to check that M is a linear map from $\mathcal{G}_2$ to the space $V(G) \times V(G)$ matrices, and it also respects products in the following sense:

$$M(x \circ y) = M(x)DM(y), \qquad M(xy) = M(x) \circ M(y)$$

(here $M(x) \circ M(y)$ denotes the Schur, or elementwise, product of these matrices). Furthermore, interchanging the labels 1 and 2 corresponds to transposition of the corresponding matrix. Clearly, $M(K_2) = B$ is the weighted adjacency matrix of G.

Now let $\mathcal{SP} \subseteq \mathcal{G}_2$ denote the space of series-parallel quantum graphs, and let $\mathcal{L}$ be the set of corresponding matrices. We want to show that the identity matrix I is in $\mathcal{L}$. Clearly $\mathcal{L}$ is a linear space that is also closed under the Schur product, the operation $(X, Y) \mapsto XDY$, and transposition. So the theorem follows if we prove the following algebraic fact.

Lemma 4.2 *Let $\mathcal{L}$ be a linear space of $n \times n$ matrices, and let D be a diagonal matrix with positive entries in the diagonal. Assume that $\mathcal{L}$ is closed under transposition, Schur product and the operation $(X, Y) \mapsto XDY$. Assume furthermore that no row is 0 in every matrix in $\mathcal{L}$, and no two rows are parallel in every matrix in $\mathcal{L}$. Then $\mathcal{L}$ contains the identity matrix.*

Proof. We start with a remark.

Claim 4.1 *Let $f : \mathbf{R} \to \mathbf{R}$ be any function and $M \in \mathcal{L}$. Then the matrix $f(M)$, obtained by applying f to every entry of M, is also in $\mathcal{L}$.*

Indeed, on the finite number of real numbers occurring as entries of M, the function f equals to some polynomial $\sum_{i=0}^{N} a_i x^i$. Then

$$f(M) = \sum_{i=0}^{N} a_i M^{(i)},$$

where $M^{(i)}$ is the Schur product of i copies of M. This shows that $f(M) \in \mathcal{L}$.

Clearly, there is a "generic" element $W \in \mathcal{L}$ such that no row or column of W is 0 and no two rows or columns of W are equal. Replacing W by $W \circ W + \varepsilon W$ with a small enough ε, we may also assume that $W \geq 0$.

We claim that all maximal entries of $W^T D W$ are on the diagonal. Indeed, suppose that $(W^T D W)_{ij}$ is a maximal entry. By Cauchy–Schwartz, we have

$$(W^T D W)_{ij} = \sum_{h \in V(G)} \alpha_h W_{hi} W_{hj} \leq \Big(\sum_{h \in V(G)} \alpha_h W_{hi}^2 \Big)^{1/2} \Big(\sum_{h \in V(G)} \alpha_h W_{hj}^2 \Big)^{1/2}$$
$$= ((W^T D W)_{ii})^{1/2}((W^T D W)_{jj})^{1/2}.$$

It follows that $(W^T D W)_{ij} = (W^T D W)_{ii} = (W^T D W)_{jj}$, and that the i-th column of W is parallel to the j-th. Since $W \geq 0$, this implies that the i-th column is equal to the j-th, and hence by the choice of W it follows that $i = j$.

Applying Claim 4.1, we can replace the maximal entries of $W^{\mathsf{T}} D W$ by 1 and all the other entries by 0, to get a nonzero diagonal 0-1 matrix $P \in \mathcal{L}$. Choose such a matrix P with maximum rank; we claim that it is the identity matrix.

Suppose not, and consider the matrix $Q = I - P$ (we don't know yet that $Q \in \mathcal{L}$). The matrix $PDP = PD = DP$ is in $\mathcal{L}$, and applying Claim 4.1 again, we get that $PD^{-1}P = PD^{-1} = D^{-1}P \in \mathcal{L}$. Hence for every matrix $M \in \mathcal{L}$,

$$QMQ = M - MP - PM + PMP$$
$$= M - MD(D^{-1}P) - (PD^{-1})DM + (PD^{-1})DMD(D^{-1}P) \in \mathcal{L}.$$

In particular, $QW^T D W Q \in \mathcal{L}$. By the same argument as above, we see that all maximal entries of $QW^T D W Q$ are on its diagonal, and so applying

15

Claim 4.1 again, we get that M contains a nonzero matrix Q' obtained from Q by changing some of its 1's to 0. Now $P + Q' \in \mathcal{L}$ is a diagonal 0-1 matrix with larger rank than P, a contradiction. $\qquad\square$

4.2 Proof of Theorem 1.5

The necessity of the conditions follows by Theorem 1.4.

To prove the sufficiency of the conditions, it suffices to prove that there exists a $q > 0$ such that $\mathrm{rk}(M(f,k)) \leq q^k$ for all $k \geq 0$, and then invoke Theorem 1.1. Note that reflection positivity is used twice: the existence of a contractor does not in itself imply an exponential bound on the rank connectivity (cf. Example 2.2). Let g_0 be a contractor for f; we show that $q = f(g_0^2)$ satisfies these conditions. Since f is multiplicative, we already know this for $k = 0$.

We may normalize f so that $f(K_1) = 1$. Let $N = \mathrm{rk}(M(f,k))$. Consider the basic idempotents $p_1, \dots, p_N$ in the algebra of k-labeled quantum graphs defined by f, and let $q_i = \frac{1}{\sqrt{f(p_i)}} p_i$. Let $q_i \otimes q_i$ denote the $(2k)$-labeled quantum graph obtained from $2k$ labeled nodes by attaching a copy of q_i at $\{1, \dots, k\}$ and another copy of q_i at $[k+1, 2k]$. Let h denote the $(2k)$-labeled quantum graph obtained from $2k$ labeled nodes by attaching a copy of g_0 at $\{i, k+i\}$ for each $i = 1, \dots, k$. Consider the quantum graph

$$ x = \sum_{i=1}^{N} q_i \otimes q_i - h. $$

By reflection positivity, we have $f(x^2) \geq 0$. But

$$ f(x^2) = \sum_{i=1}^{N}\sum_{i=1}^{N} \langle q_i \otimes q_i, q_i \otimes q_i \rangle - 2\sum_{i=1}^{N} \langle q_i \otimes q_i, h \rangle + \langle h, h \rangle. $$

Here by the fact that $q_i q_j = 0$ if $i \neq j$, we have

$$ \sum_{i=1}^{N}\sum_{j=1}^{N} \langle q_i \otimes q_i, q_j \otimes q_j \rangle = \sum_{i=1}^{N} \langle q_i, q_i \rangle^2 = N. $$

Furthermore, by the definition of g_0 and h, we have

$$ \sum_{i=1}^{N} \langle q_i \otimes q_i, h \rangle = \sum_{i=1}^{N} \langle q_i, q_i \rangle = N. $$

Finally, by the definition of h and the multiplicativity of f, we have

$$\langle h, h \rangle = f(g_0^2)^k.$$

Thus $f(x^2) \geq 0$ implies that $N \leq f(g_0^2)^k$, which completes the proof.

4.3 Proof of Theorem 1.6

Set $t = t(.,W)$. Let $F : [0,1]^2 \to [0,1]$ be any integrable function. We define

$$\|F\|_x = \int_0^1 |F(x,y)|\, dy.$$

Then

$$\|F\| = \int_0^1 \|F\|_x\, dx \tag{7}$$

is the usual ℓ_1-norm of F.

Lemma 4.3 *Let $U, W : [0,1]^2 \to [0,1]$ be two symmetric functions and let F be a 2-labeled graph with m edges in which the labeled nodes are independent. Let d_i denote the degree of i in F. Then for every $x, y \in [0,1]$,*

$$|t_{xy}(F,U) - t_{xy}(F,W)| \leq d_1\|U-W\|_x + d_2\|U-W\|_y + (m - d_1 - d_2)\|U-W\|.$$

Proof. Let $V(F) = \{1,\ldots,n\}$ and $E(F) = \{e_1,\ldots,e_m\}$, where $e_t = i_t j_t$, $i_t < j_t$. Then

$$t_{x_1 x_2}(F,U) - t_{x_1 x_2}(F,W)$$
$$= \int_{[0,1]^{n-2}} \left(\prod_{ij \in E(F)} W(x_i, x_j) - \prod_{ij \in E(F)} U(x_i, x_j) \right) dx_3 \ldots dx_n.$$

We can write

$$\prod_{ij \in E(F)} W(x_i, x_j) - \prod_{ij \in E(F)} U(x_i, x_j) = \sum_{t=1}^m X_t(x),$$

where

$$X_t(x) = \left(\prod_{s=1}^{t-1} W(x_{i_s}, x_{j_s}) \right) \left(\prod_{s=t+1}^m U(x_{i_s}, x_{j_s}) \right) (W(x_{i_t}, x_{j_t}) - U(x_{i_t}, x_{j_t})).$$

17

Consider the integral of a given term:

$$\left| \int\limits_{[0,1]^{n-2}} X_t(x) \, dx_3 \ldots dx_n \right| \leq \int\limits_{[0,1]^n} |X_t(x)| \, dx_3 \ldots dx_n$$

$$\leq \int\limits_{[0,1]^{n-2}} |W(x_{i_t}, x_{j_t}) - U(x_{i_t}, x_{j_t})| \, dx_3 \ldots dx_n.$$

If $i_t = 1$, then this integral is just $\|W - U\|_{x_1}$; if $i_t = 2$, then it is $\|W - U\|_{x_1}$; if $i_t \geq 3$, then it is $\|W - U\|$. (Note that $i_t = 1, j_t = 2$ does not occur by hypothesis.) The first possibility occurs d_1 times, the second d_2 times. This proves the Lemma. $\qquad\qquad\square$

Remark. In [4] a version of this lemma was proved (for unlabeled graphs) where the ℓ_1 norm was replaced by the smaller "rectangle norm". Such a sharper version could be proved here as well (but we don't need it).

Applying lemma 4.3 to all simple 2-labeled graphs occurring in a quantum graph, we get

Corollary 4.4 *Let $U, W : [0,1]^2 \to [0,1]$ be two symmetric functions and let g be any simple 2-labeled quantum graph. Then there exists a constant $c = c(g)$ depending only on g such that for every $x, y \in [0,1]$,*

$$|t_{xy}(g, U) - t_{xy}(g, W)| \leq c_g \max(\|U - W\|_x, \|U - W\|_y, \|U - W\|).$$

Now we return to the proof of of the Theorem. Let g be a simple 2-labeled quantum graph and assume that $g \equiv 0 \pmod{t}$. Then in particular

$$t(gg, W) = \int\limits_0^1 \int\limits_0^1 t_{xy}(g, W)^2 \, dx \, dy = 0,$$

and hence

$$t_{xy}(g, W) = 0 \qquad\qquad (8)$$

for almost all $x, y \in [0,1]$. Let $C \subset [0,1]$ denote the set where this does not hold.

Next we show that

$$t_{xx}(g, W) = 0 \qquad\qquad (9)$$

for almost all $x \in [0,1]$.

Suppose (9) is false; then there is an $\varepsilon > 0$ and a set $A \subseteq [0,1]$ with $\lambda(A) = \varepsilon$ such that (say) $t_{xx}(g, W) > \varepsilon$ for all $x \in A$. Let U be a continuous function such that

$$\|U - W\| < \frac{\varepsilon^2}{9c_g}.$$

From (7) it follows that the set

$$B = \{x \in [0,1] : \ \|U - W\|_x > \frac{\varepsilon}{3c_g}\}$$

has measure less than $\varepsilon/3$. For $x \in [0,1]$, let

$$C_x = \{y \in [0,1] : \ (x, y) \in C\},$$

and let D be the set of points $x \in [0,1]$ for which the set C_x does not have measure 0. Clearly, D has measure 0. Hence $A \backslash B \backslash D$ has positive measure.

Let x be a point with density 1 of the set $A \backslash B \backslash D$. Choose any sequence $y_n \in A \backslash B \backslash C_x$ such that $y_n \to x$ (such a sequence exists since C_x has measure 0 by the choice of x). Then we have by Corollary 4.4

$$|t_{xy_n}(g, U) - t_{xy_n}(g, W)| < c_g \max(\frac{\varepsilon}{3c_g}, \frac{\varepsilon}{3c_g}, \frac{\varepsilon^2}{9c_g}) = \frac{\varepsilon}{3}, \tag{10}$$

since $x, y_n \notin B$, and similarly,

$$|t_{xx}(g, U) - t_{xx}(g, W)| < \frac{\varepsilon}{3}. \tag{11}$$

Here $t_{xy_n}(g, W) = 0$ since $y_n \notin C_x$, and $t_{xx}(g, W) > \varepsilon$, since $x \in A$. Thus (10) and (11) imply that

$$|t_{xy_n}(g, U) - t_{xx}(g, U)| > \varepsilon - 2\frac{\varepsilon}{3} > \frac{\varepsilon}{3},$$

which is a contradiction, since U is continuous, and therefore

$$t_{xy_n}(g, U) \to t_{xx}(g, U).$$

This contradiction proves (9).

From here, the proof of Theorem 1.6 is easy. Trivially

$$t_x(g', W) = t_{xx}(G, W),$$

and so (9) implies that for every 1-labeled quantum graph h

$$t(g'h, W) = \int_0^1 t_x(g'h, W) \, dx = \int_0^1 t_x(g', W) t_x(h, W) \, dx = 0.$$

19

To prove the second assertion of the theorem, it suffices to note that if t had a contractor, then it would have a representation in the form of $t = t(., H)$ with some finite weighted graph H by Theorem 1.5. In other words, we would have a stepfunction W' such that $t(F, W) = t(F, W')$ for every finite graph F. By the results of [1], this implies that W is a stepfunction (up to set of measure 0).

References

[1] C. Borgs, J. Chayes and L. Lovász: Unique limits of dense graph sequences (manuscript)

[2] M. Freedman, L. Lovász, A. Schrijver: Reflection positivity, rank connectivity, and homomorphism of graphs (MSR Tech Report # MSR-TR-2004-41)

   ftp://ftp.research.microsoft.com/pub/tr/TR-2004-41.pdf

[3] L. Lovász: The rank of connection matrices and the dimension of graph algebras, *Eur. J. Combinatorics* (to appear).

[4] L. Lovász and B. Szegedy: Limits of dense graph sequences (MSR Tech Report # MSR-TR-2004-79)

   ftp://ftp.research.microsoft.com/pub/tr/TR-2004-79.pdf

[5] L. Lovász and V.T. Sós: Generalized quasirandom graphs (manuscript).

[6] M. Freedman, L. Lovász, D. Welsh (unpublished)