Notation

We follow the basic graph theory notation from the book of Bondy and Murty [1]. For a graph G, we denote the set of vertices of G by $V(G)$ and the set of edges of G by $E(G) \subseteq \binom{V(G)}{2}$. We call the cardinality of $V(H)$ the *order* of H and denote it by $v(H)$, and the cardinality of $E(H)$ the *size* of H and denote it by $e(H)$.

For a graph H, we denote by $\overline{H}$ the complement of H, i.e., the graph with the vertex-set $V(H)$ and the edge-set $\binom{V(H)}{2} \setminus E(H)$. For a subset of vertices $S \subseteq V(H)$, we denote by $H[S]$ the induced subgraph, i.e., the graph with the vertex-set S and the edge-set $\{e \in E(H) : e \subseteq S\}$.

1 Flag Algebras

We follow the approach of Razborov [4] and introduce the framework of flag algebras for the graphs. Exactly the same scheme can also be used to setup flag algebras for the oriented graphs, the ℓ-uniform hypergraphs (ℓ-graphs), the permutations, and many others. In fact, Razborov introduced in [4] the framework for an arbitrary universal first-order logic theory without constants or function symbols. We decided to present in this section the flag algebra setup for the particular instance of graphs rather than in the general setting, since it might be easier to understand the ideas of the framework in this way.

1.1 Flag algebra setting for graphs

The central notions we are going to introduce are an algebra $\mathcal{A}$ and algebras $\mathcal{A}^\sigma$, where σ is a fixed graph with a fixed labelling of its vertex set. In order to precisely describe the algebras $\mathcal{A}$ and $\mathcal{A}^\sigma$ on formal linear combinations of graphs, we first need to introduce some additional notation. Let $\mathcal{F}$ be the set of all finite non-isomorphic graphs. Next, for every $\ell \in \mathbb{N}$, let $\mathcal{F}_\ell \subset \mathcal{F}$ be the set of all graphs of order ℓ. For convenience, we fix an arbitrary ordering on the elements the set $\mathcal{F}_\ell$ for every $\ell \in \mathbb{N}$, i.e., we always assume that $\mathcal{F}_\ell = \{F_1, F_2, \ldots, F_{|\mathcal{F}_\ell|}\}$.

For $H \in \mathcal{F}_\ell$ and $H' \in \mathcal{F}_{\ell'}$, we define $p(H, H')$ to be the probability that a randomly chosen subset of ℓ vertices in H' induces a subgraph isomorphic to H. Note that $p(H, H') = 0$ if $\ell' < \ell$. Let $\mathbb{R}\mathcal{F}$ be the set of all formal linear combinations of elements of $\mathcal{F}$ with real coefficients. Furthermore, let $\mathcal{K}$ be the linear subspace of $\mathbb{R}\mathcal{F}$ generated by all the linear combinations of the form

$$H - \sum_{H' \in \mathcal{F}_{v(H)+1}} p(H, H') \cdot H'.$$

Two examples of such linear combinations are depicted in Figure 1. Finally, we set $\mathcal{A}$ to be the space $\mathbb{R}\mathcal{F}$ factored by $\mathcal{K}$, and the element corresponding to $\mathcal{K}$ in $\mathcal{A}$ to be the zero element of $\mathcal{A}$.

The space $\mathcal{A}$ comes with a natural definition of an addition and a multiplication by a real number. We now introduce the notion of a product of two elements from $\mathcal{A}$. We start with the definition for the elements of $\mathcal{F}$. For $H_1, H_2 \in \mathcal{F}$, and $H \in \mathcal{F}_{v(H_1)+v(H_2)}$, we define $p(H_1, H_2; H)$ to be the probability that a randomly chosen subset of $V(H)$ of size $v(H_1)$ and its complement induce in H subgraphs isomorphic to H_1 and H_2,

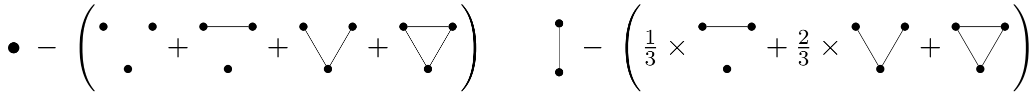

Figure 1: Two examples of linear combinations used in generating $\mathcal{K}$.

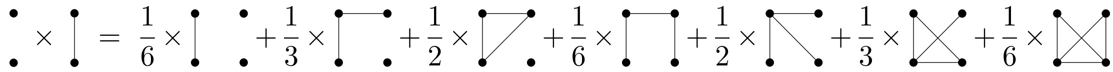

Figure 2: An example of a product in the algebra $\mathcal{A}$.

respectively. We set

$$H_1 \times H_2 := \sum_{H \in \mathcal{F}_{v(H_1)+v(H_2)}} p(H_1, H_2; H) \cdot H.$$

See Figure 2 for an example of a product with $H_1 =$ ⁚ and $H_2 =$ ⁞. The multiplication on $\mathcal{F}$ has a unique linear extension to $\mathbb{R}\mathcal{F}$, which yields a well-defined multiplication also in the factor algebra $\mathcal{A}$. A formal proof of this is given in [4, Lemma 2.4]. Observe that the one-vertex graph $\bullet \in \mathcal{F}$ is, modulo $\mathcal{K}$, the neutral element of the product in $\mathcal{A}$.

Let us now move to the definition of the algebra $\mathcal{A}^\sigma$, where σ is a fixed finite graph with a fixed labelling of its vertices. The labelled graph σ is usually called a *type*. We follow the same lines as in the definition of $\mathcal{A}$. Let $\mathcal{F}^\sigma$ be the set of all finite graphs H with a fixed *embedding* of σ, i.e., an injective mapping θ from $V(\sigma)$ to $V(H)$ such that θ is an isomorphism between σ and $H[\mathrm{Im}(\theta)]$. The elements of $\mathcal{F}^\sigma$ are usually called *σ-flags*, the subgraph induced by $\mathrm{Im}(\theta)$ is called the *root* of a σ-flag, and the vertices $\mathrm{Im}(\theta)$ are called the *rooted* or the *labelled* vertices. The vertices that are not rooted are called the *non-rooted* or the *non-labelled* vertices. For every $\ell \in \mathbb{N}$, we define $\mathcal{F}^\sigma_\ell \subset \mathcal{F}^\sigma$ to be the set of all ℓ-vertex σ-flags from $\mathcal{F}^\sigma$. Also, for each type σ and each integer ℓ, we fix an arbitrary ordering on the elements of the set $\mathcal{F}^\sigma_\ell$.

In the analogy to the case of $\mathcal{A}$, for two σ-flags $H \in \mathcal{F}^\sigma$ and $H' \in \mathcal{F}^\sigma$ with the embeddings of σ given by θ and θ', respectively, we set $p(H, H')$ to be the probability that a randomly chosen subset of $v(H) - v(\sigma)$ vertices in $V(H') \setminus \theta'(V(\sigma))$ together with $\theta'(V(\sigma))$ induces a subgraph that is isomorphic to H through an isomorphism f that preserves the embedding of σ. In other words, the isomorphism f has to satisfy $f(\theta') = \theta$. Let $\mathbb{R}\mathcal{F}^\sigma$ be the set of all formal linear combinations of elements of $\mathcal{F}^\sigma$ with real coefficients, and let $\mathcal{K}^\sigma$ be the linear subspace of $\mathbb{R}\mathcal{F}^\sigma$ generated by all the linear combinations of the form

$$H - \sum_{H' \in \mathcal{F}^\sigma_{v(H)+1}} p(H, H') \cdot H'.$$

See Figure 3 for two examples of such linear combinations in the case σ is the one-vertex type. We define $\mathcal{A}^\sigma$ to be $\mathbb{R}\mathcal{F}^\sigma$ factored by $\mathcal{K}^\sigma$ and, analogously to the case for the algebra $\mathcal{A}$, we let the element corresponding to $\mathcal{K}^\sigma$ to be the zero element of $\mathcal{A}^\sigma$.

We now define the product of two elements from $\mathcal{F}^\sigma$. Let $H_1, H_2 \in \mathcal{F}^\sigma$ and $H \in \mathcal{F}^\sigma_{v(H_1)+v(H_2)-v(\sigma)}$ be σ-flags, and θ be the fixed embedding of σ in H. Similarly to the

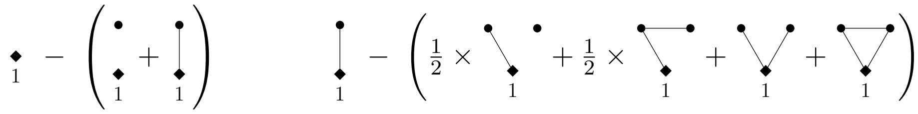

Figure 3: Two examples of linear combinations used in generating $\mathcal{K}^\sigma$, where σ is the one-vertex type.

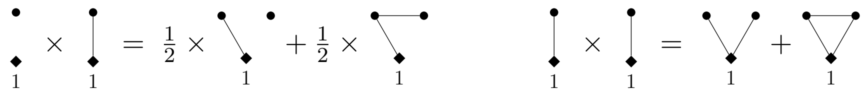

Figure 4: Two examples of a product in the algebra $\mathcal{A}^\sigma$, where σ is the one-vertex type.

definition of the multiplication for $\mathcal{A}$, we define $p(H_1, H_2; H)$ to be the probability that a randomly chosen subset of $V(H) \setminus \theta(V(\sigma))$ of size $v(H_1) - v(\sigma)$ and its complement in $V(H) \setminus \theta(V(\sigma))$ of size $v(H_2) - v(\sigma)$, extend $\theta(V(\sigma))$ in H to subgraphs isomorphic to H_1 and H_2, respectively. Again, by isomorphic we mean that there is an isomorphism that preserves the fixed embedding of σ. We set

$$H_1 \times H_2 := \sum_{H \in \mathcal{F}^\sigma_{v(H_1)+v(H_2)-v(\sigma)}} p(H_1, H_2; H) \cdot H.$$

Two examples of a product in $\mathcal{A}^\sigma$ for σ being the one-vertex type are depicted in Figure 4. The definition of the product for the elements of $\mathcal{F}^\sigma$ naturally extends to $\mathcal{A}^\sigma$. It follows that the unique σ-flag of size $v(\sigma)$ represents, modulo $\mathcal{K}^\sigma$, the neutral element of the product in $\mathcal{A}^\sigma$.

Now consider an infinite sequence $(G_n)_{n\in\mathbb{N}}$ of graphs with increasing orders. We say that the sequence $(G_n)_{n\in\mathbb{N}}$ is *convergent* if the probabilities $p(H, G_n)$ converge for every $H \in \mathcal{F}$. A standard compactness argument (e.g., using Tychonoff's theorem [5]) yields that every such infinite sequence has a convergent subsequence.

Fix a convergent sequence $(G_n)_{n\in\mathbb{N}}$ of graphs with increasing orders. For every $H \in \mathcal{F}$, we set $\phi(H) = \lim_{n\to\infty} p(H, G_n)$, and we then linearly extend ϕ to $\mathcal{A}$. We usually refer to the mapping ϕ as to the *limit* of the sequence. The obtained mapping ϕ is a homomorphism from $\mathcal{A}$ to $\mathbb{R}$, see [4, Theorem 3.3a]. Moreover, for every $H \in \mathcal{F}$, it holds $\phi(H) \geq 0$. Let $\mathrm{Hom}^+(\mathcal{A}, \mathbb{R})$ be the set of all such homomorphisms, i.e., the set of all homomorphisms ψ from the algebra $\mathcal{A}$ to $\mathbb{R}$ such that $\psi(H) \geq 0$ for every $H \in \mathcal{F}$. It is interesting to see that this set is exactly the set of all the limits of convergent sequences of graphs [4, Theorem 3.3b].

Let $(G_n)_{n\in\mathbb{N}}$ be a convergent sequence of graphs and $\phi \in \mathrm{Hom}^+(\mathcal{A}, \mathbb{R})$ its limit. For a type σ and an embedding θ of σ in G_n, we define G_n^θ to be the graph rooted on the copy of σ that corresponds to θ. For every $n \in \mathbb{N}$ and $H^\sigma \in \mathcal{F}^\sigma$, we define $p_n^\theta(H^\sigma) = p(H^\sigma, G_n^\sigma)$. Picking θ at random gives rise to a probability distribution $\mathbf{P}^\sigma_\mathbf{n}$ on mappings from $\mathcal{A}^\sigma$ to $\mathbb{R}$, for every $n \in \mathbb{N}$. Since $p(H, G_n)$ converges (as n tends to infinity) for every $H \in \mathcal{F}$, the sequence of these probability distributions on mappings from $\mathcal{A}^\sigma$ to $\mathbb{R}$ weakly converges to a Borel probability measure on $\mathrm{Hom}^+(\mathcal{A}^\sigma, \mathbb{R})$, see [4, Theorems 3.12 and 3.13]. We denote the limit probability distribution by $\mathbf{P}^\sigma$. In fact, for any σ such that $\phi(\sigma) > 0$, the homomorphism ϕ itself fully determines the probability distribution $\mathbf{P}^\sigma$ [4, Theorem 3.5]. Furthermore, any mapping ϕ^σ from the support of the distribution $\mathbf{P}^\sigma$ is in fact a homomorphism from $\mathcal{A}^\sigma$ to $\mathbb{R}$ such that $\phi^\sigma(H^\sigma) \geq 0$ for

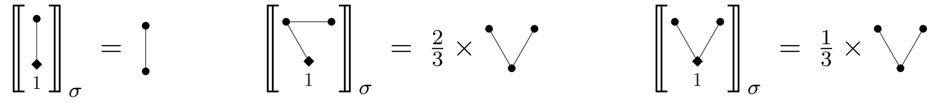

Figure 5: Three examples of applying the averaging operator $[\![\cdot]\!]_\sigma$, where σ denotes the one-vertex type.

all $H^\sigma \in \mathcal{F}^\sigma$ [4, Proof of Theorem 3.5].

The last notion we introduce is the *averaging* (or downward) operator $[\![\cdot]\!]_\sigma : \mathcal{A}^\sigma \to \mathcal{A}$. It is the linear operator defined on the elements of $H \in \mathcal{F}^\sigma$ by

$$[\![H]\!]_\sigma := p_H^\sigma \cdot H^\emptyset,$$

where $H^\emptyset$ is the (unlabelled) graph from $\mathcal{F}$ corresponding to H after unlabelling all its vertices, and p_H^σ is the probability that a random injective mapping from $V(\sigma)$ to $V(H^\emptyset)$ is an embedding of σ in $H^\emptyset$ yielding a σ-flag isomorphic to H. See Figure 5 for three examples of applying the averaging operator $[\![\cdot]\!]_\sigma$, where σ is the one-vertex type.

The key relation between ϕ and ϕ^σ is the following

$$\forall H^\sigma \in \mathcal{A}^\sigma, \quad \phi\left([\![H^\sigma]\!]_\sigma\right) = \phi\left([\![\sigma]\!]_\sigma\right) \cdot \int \phi^\sigma(H^\sigma), \tag{1}$$

where the integration is with respect to the probability measure given by the random distribution $\mathbf{P}^\sigma$ on ϕ^σ. Note that

$$\phi\left([\![\sigma]\!]_\sigma\right) = \frac{\left|\mathrm{Aut}\left(\sigma^\emptyset\right)\right|}{v\left(\sigma^\emptyset\right)!} \cdot \phi\left(\sigma^\emptyset\right).$$

Vaguely speaking, the relation 1 corresponds to the conditional probability formula $\mathbb{P}\big[A \cap B\big] = \mathbb{P}\big[B\big] \cdot \mathbb{P}\big[A \mid B\big]$, where B is the event that a random injective mapping θ is an embedding of σ, and A is the event that a random subset of $v(H) - v(\sigma)$ vertices extends θ to the σ-flag H^σ. A formal proof is given in [4, Lemma 3.11]. The relation (1) implies that if $\phi^\sigma(A^\sigma) \geq 0$ almost surely for some $A^\sigma \in \mathcal{A}^\sigma$, then $\phi\left([\![A^\sigma]\!]_\sigma\right) \geq 0$. In particular, for every homomorphism $\phi \in \mathrm{Hom}^+(\mathcal{A}, \mathbb{R})$ and every linear combination $A^\sigma \in \mathcal{A}^\sigma$ it holds

$$\phi\left(\left[\!\!\left[(A^\sigma)^2\right]\!\!\right]_\sigma\right) \geq 0. \tag{2}$$

We note that a stronger variant of (2) can be proven using Cauchy-Schwarz's inequality. Specifically, [4, Theorem 3.14] states that

$$\forall \phi \in \mathrm{Hom}^+(\mathcal{A}, \mathbb{R}), \forall A^\sigma, B^\sigma \in \mathcal{A}^\sigma,\ \phi\left(\left[\!\!\left[(A^\sigma)^2\right]\!\!\right]_\sigma \times \left[\!\!\left[(B^\sigma)^2\right]\!\!\right]_\sigma\right) \geq \phi\left([\![A^\sigma \times B^\sigma]\!]_\sigma\right)^2.$$

Let σ be a type, $A^\sigma \in \mathcal{A}^\sigma$, and m the minimum integer such that

$$A^\sigma = \sum_{F^\sigma \in \mathcal{F}_m^\sigma} \alpha_{F^\sigma} \cdot F_i^\sigma.$$

We say that $\ell := 2m - v(\sigma)$ is the *order* of the expression $\left[\!\!\left[(A^\sigma)^2\right]\!\!\right]_\sigma$. It follows that

$$\left[\!\!\left[(A^\sigma)^2\right]\!\!\right]_\sigma = \sum_{F \in \mathcal{F}_\ell} \alpha_F \cdot F_i.$$

Since the operator $[\![\cdot]\!]_\sigma$ is linear, it immediately follows that for every $\phi \in \mathrm{Hom}^+(\mathcal{A}, \mathbb{R})$, $A_1^\sigma, A_2^\sigma, \dots A_K^\sigma \in \mathcal{A}^\sigma$, and non-negative reals $\alpha_1, \alpha_2, \dots, \alpha_K$ we have

$$\phi\left(\left[\!\!\left[\alpha_1 \cdot (A_1^\sigma)^2 + \alpha_2 \cdot (A_2^\sigma)^2 + \dots + \alpha_K \cdot (A_K^\sigma)^2\right]\!\!\right]_\sigma\right) \geq 0.$$

Hence for every finite set $S \subseteq \mathcal{F}^\sigma$ and every real symmetric positive semidefinite matrix M of size $|S| \times |S|$ it holds that

$$\phi\left([\![x_S^T M x_S]\!]_\sigma\right) \geq 0,$$

where x_S is the $|S|$-dimensional vector from whose i-th coordinate is equal to the i-th element of S. Note that we used the fact that every real symmetric positive semidefinite matrix M of size $s \times s$ can be written as a sum of squares. In other words, there exists an integer $d \leq s$ such that

$$M = \sum_{j \in [d]} \lambda_j \cdot v_j \times (v_j)^T ,$$

where vectors v_i, for $i \in [d]$, form an orthonormal eigen-basis of M and λ_i are the corresponding (always non-negative) eigenvalues. On the other hand, for every set of d' non-negative reals λ'_j and vectors $v'_j \in \mathbb{R}^s$, the matrix

$$M' := \sum_{j \in [d']} \lambda'_j \cdot v'_j \times \left(v'_j\right)^T$$

is a symmetric positive semidefinite matrix of size $s \times s$.

1.2 Semidefinite method

The heart of the semidefinite method are inequalities of the form (2). These inequalities turned out to be very useful while dealing with problems from extremal combinatorics. Before moving to the description of the method itself, let us illustrate this on a small example. Consider the following inequality, where σ is the one-vertex type:

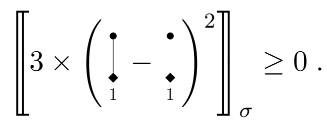

(3)

The definition of the product on A^σ implies that the left-hand side of (3) is equal to

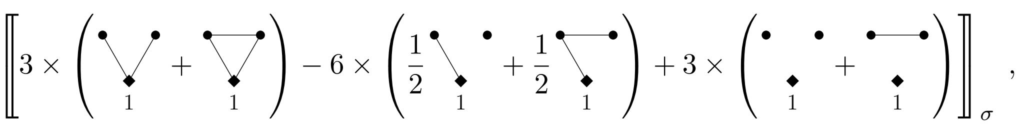

and an application of the averaging operator yields that

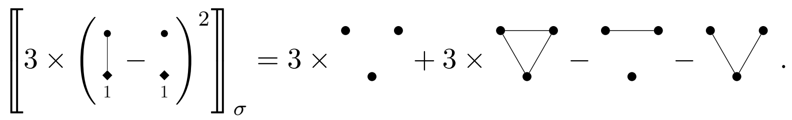

Therefore, every $\phi \in \mathrm{Hom}^+(\mathcal{A}, \mathbb{R})$ satisfies

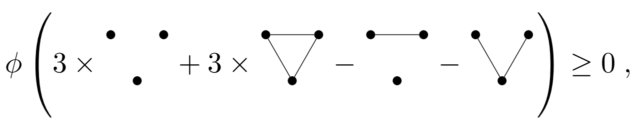

and hence also

$$\phi\left(4\times$$ 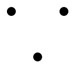 $$+\,4\times$$ 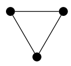 $$\right)\geq\phi\left(\right.$$ 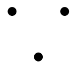 $$+$$ 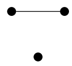 $$+$$ 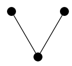 $$+$$ 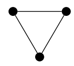 $$\left.\right).$$

Since the right-hand side of the last inequality is equal to one, we conclude that

$$\phi\left(\right.$$ 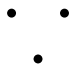 $$+$$ 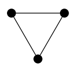 $$\left.\right)\geq\frac{1}{4} \tag{4}$$

for every $\phi \in \mathrm{Hom}^+(\mathcal{A}, \mathbb{R})$. This is a well-known inequality due to Goodman [2]. Note that the inequality (4) is best possible. This can be seen, for example, by considering the limit of the sequence of Erdős-Renyi random graphs $G_{n,1/2}$ with increasing orders (the sequence is convergent with probability 1). Another example where the inequality (4) is tight is the sequence of complete balanced bipartite graphs with increasing orders (it is straightforward to check that this sequence is convergent).

Now consider a general linear combination $A \in \mathcal{A}$. One of the fundamental problems in extremal combinatorics is to determine the smallest value of $\phi(A)$ over all $\phi \in \mathrm{Hom}^+(\mathcal{A}, \mathbb{R})$. The semidefinite method is a tool from the flag algebra framework that systematically searches for inequalities of the form (2), like the inequality (3) in the case when $A =$ 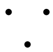 $+$ 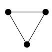, in order to find a lower bound on

$$\min_{\phi \in \mathrm{Hom}^+(\mathcal{A},\mathbb{R})} \phi(A). \tag{5}$$

Note that since $\mathrm{Hom}^+(\mathcal{A}, \mathbb{R})$ is compact, such a minimum exists for every $A \in \mathcal{A}$.

The semidefinite method works as follows. First, fix an upper bound ℓ on the order of flags in all the terms of linear inequalities we are going to consider, including also the terms of the objective function A. Without loss of generality, $A = \sum_{F \in \mathcal{F}_\ell} \alpha_F \cdot F$. Next, fix an arbitrary list of types $\sigma_1, \ldots, \sigma_K$ of order at most ℓ. Recall our aim is to find a lower bound on (5). The semidefinite method finds a way how to express A in the algebra $\mathcal{A}$ as follows:

$$A = \underbrace{\left(\sum_{k\in[K]}\sum_{j\in[J_k]} \beta_j^k \cdot \left[\!\!\left[\left(A_j^{\sigma_k}\right)^2\right]\!\!\right]_{\sigma_k}\right)}_{R} + \underbrace{\left(\sum_{F\in\mathcal{F}_\ell} \beta_F \cdot F\right)}_{S} + \underbrace{\left(c \cdot \sum_{F\in\mathcal{F}_\ell} F\right)}_{T}, \tag{6}$$

where

- $J_1, \ldots J_K$ are non-negative integers,
- $A_j^{\sigma_1} \in \mathcal{A}^{\sigma_1}$ so that the order of $\left[\!\!\left[\left(A_j^{\sigma_1}\right)^2\right]\!\!\right]_{\sigma_1}$ is at most ℓ for every $j \in [J_1]$,

$$\vdots$$

- $A_j^{\sigma_K} \in \mathcal{A}^{\sigma_K}$ so that the order of $\left[\!\!\left[\left(A_j^{\sigma_K}\right)^2\right]\!\!\right]_{\sigma_K}$ is at most ℓ for every $j \in [J_K]$,

- $\beta^k_j \geq 0$ for every $k \in [K]$ and $j \in [J_k]$,
- $\beta_F \geq 0$ for every $F \in \mathcal{F}_\ell$, and
- $c \in \mathbb{R}$.

Since $\phi(R) \geq 0$, $\phi(S) \geq 0$, and $\phi(T) = c$ for all $\phi \in \mathrm{Hom}^+(\mathcal{A}, \mathbb{R})$, we conclude that $\phi(A) \geq c$. Note that R is a positive linear combination of inequalities (2) of order at most ℓ, hence $R = \sum_{F \in \mathcal{F}_\ell} r_F \cdot F$ for some choice of reals r_F.

For a fixed choice of the parameters ℓ and $\sigma_1, \ldots, \sigma_K$, finding such an expression of A can be formulated as a semidefinite program. Note that all the expressions A, R, S and T can be written as linear combinations of the elements from $\mathcal{F}_\ell$, i.e., they can be viewed as vectors in $\mathbb{R}^{|\mathcal{F}_\ell|}$. Furthermore, the bound obtained by the semidefinite method is "best possible" in the following sense. Let c_0 be the obtained bound. For every expression of A as a linear combination of the form (6), the coefficient c in this combination is at most c_0. Note that there might be (and often there are) different combinations that yield the bound c_0.

Let us now describe the corresponding semidefinite program in more detail. Fix one of the types $\sigma_k \in \{\sigma_1, \ldots, \sigma_K\}$. Since the semidefinite method uses inequalities $\left[\!\left[(A^{\sigma_k})^2 \right]\!\right]_{\sigma_k}$ of order at most ℓ, it follows that

$$A^{\sigma_k} = \sum_{F_i \in \mathcal{F}^{\sigma_k}_{m(k)}} \alpha_i \cdot F_i$$

for some integer $m(k)$ such that $m(k) \leq \frac{\ell + v(\sigma_k)}{2}$. Without loss of generality, $m(k) = \left\lfloor \frac{\ell + v(\sigma_k)}{2} \right\rfloor$. Therefore,

$$R = \sum_{k \in [K]} \sum_{j \in [J_k]} \beta^k_j \cdot \left[\!\!\left[\left(\sum_{F_i \in \mathcal{F}^{\sigma_k}_{m(k)}} \alpha_{k,j,i} \cdot F_i \right)^2 \right]\!\!\right]_{\sigma_k} = \sum_{k \in [K]} \left[\!\left[x^T_{\sigma_k} M_{\sigma_k} x_{\sigma_k} \right]\!\right]_{\sigma_k}, \qquad (7)$$

where

- each vector x_{σ_k} is the $|\mathcal{F}^{\sigma_k}_{m(k)}|$-dimensional vector whose i-th coordinate is equal to the i-th element of $\mathcal{F}^{\sigma_k}_{m(k)}$, and
- each matrix M_{σ_k} is equal to

$$\sum_{j \in [J_k]} \beta^k_j \cdot \left(\alpha_{k,j,1}, \alpha_{k,j,2}, \ldots, \alpha_{k,j,\left|\mathcal{F}^{\sigma_k}_{m(k)}\right|} \right)^T \times \left(\alpha_{k,j,1}, \alpha_{k,j,2}, \ldots, \alpha_{k,j,\left|\mathcal{F}^{\sigma_k}_{m(k)}\right|} \right).$$

Note that the matrices $M_{\sigma_1}, \ldots, M_{\sigma_K}$ are symmetric positive semidefinite matrices with real entries.

Now recall that $R = \sum_{F \in \mathcal{F}_\ell} r_F \cdot F$. The equation (7) implies that all the coefficients r_F depends only on the entries of the matrices $M_{\sigma_1}, \ldots, M_{\sigma_K}$. For a given set of matrices $M_{\sigma_1}, \ldots, M_{\sigma_K}$, we write $r_F(M_{\sigma_1}, \ldots, M_{\sigma_K})$ to denote the coefficient in front of F in R.

Using this notation, the semidefinite program for the objective value $A = \sum_{F \in \mathcal{F}_\ell} \alpha_F \cdot F$ can be written as

$$\begin{aligned} \underset{c \in \mathbb{R}}{\text{maximize}} \quad & c \\ \text{subject to} \quad & \alpha_F \geq c + r_F(M_{\sigma_1}, \dots, M_{\sigma_K}) \quad \forall F \in \mathcal{F}_\ell, \\ & M_{\sigma_1} \succeq 0, \\ & \vdots \\ & M_{\sigma_K} \succeq 0, \end{aligned} \tag{8}$$

where the constraints $M_{\sigma_k} \succeq 0$, for $k \in [K]$, denote that the matrices M_{σ_k} are positive semidefinite.

Let us now focus on the dual program of the semidefinite program (8). We start with introducing some additional notation. For a homomorphism $\phi \in \mathrm{Hom}^+(\mathcal{A}, \mathbb{R})$ and an integer ℓ, the *local density ℓ-profile* of ϕ is the vector

$$\phi_{\restriction \ell} := \left(\phi(F_1), \phi(F_2), \dots, \phi(F_{|\mathcal{F}_\ell|})\right).$$

We denote the i-th coordinate of $\phi_{\restriction \ell}$ by $\phi_{\restriction \ell}(F_i)$. Furthermore, for $A = \sum_{F \in \mathcal{F}_\ell} \alpha_F \cdot F$, where α_F are arbitrary fixed reals, we define

$$\phi_{\restriction \ell}(A) := \sum_{F \in \mathcal{F}_\ell} \alpha_F \cdot \phi_{\restriction \ell}(F).$$

With a slight abuse of notation, we use this notion also for an arbitrary vector $z \in \mathbb{R}^{|\mathcal{F}_\ell|}$, i.e., we write $z(F_i)$ for the i-th coordinate of z, and we use $z(A)$ to denote the value of $\sum_{F \in \mathcal{F}_\ell} \alpha_F \cdot z(F)$.

Let $\mathcal{P}_{\mathcal{F}_\ell} := \left\{\phi_{\restriction \ell} : \phi \in \mathrm{Hom}^+(\mathcal{A}, \mathbb{R})\right\}$ be the set of all local density ℓ-profiles. Note that $\mathcal{P}_{\mathcal{F}_\ell} \subseteq [0,1]^{|\mathcal{F}_\ell|}$. For a combination $A = \sum_{F \in \mathcal{F}_\ell} \alpha_F \cdot F$, it follows from the definitions that the value of (5) is equal to the minimum value of $\phi_{\restriction \ell}(A)$, where the minimum is taken over all $\phi_{\restriction \ell} \in \mathcal{P}_{\mathcal{F}_\ell}$.

Now fix K types $\sigma_1, \dots, \sigma_K$ of order at most ℓ. Let $\mathcal{S}_{\mathcal{F}_\ell}$ be the set of vectors $z \in \mathbb{R}^{|\mathcal{F}_\ell|}$ that satisfy

- all the linear inequalities of the form $z\left(\left[\!\left[(A^{\sigma_k})^2\right]\!\right]_{\sigma_k}\right) \geq 0$, where $k \in [K]$, $A^{\sigma_k} \in \mathcal{A}^{\sigma_k}$, and the order of the expression $\left[\!\left[(A^{\sigma_k})^2\right]\!\right]_{\sigma_k}$ is at most ℓ,
- the non-negative inequalities $z(F) \geq 0$ for every $F \in \mathcal{F}_\ell$, and
- the equation $z\left(\sum_{F \in \mathcal{F}_\ell} F\right) = 1$.

It immediately follows that $\mathcal{P}_{\mathcal{F}_\ell} \subseteq \mathcal{S}_{\mathcal{F}_\ell}$. It also follows that the set $\mathcal{S}_{\mathcal{F}_\ell}$ is a convex set.

Recall that the aim of the semidefinite method is to find the minimum of $\phi(A)$, where $\phi \in \mathrm{Hom}^+(\mathcal{A}, \mathbb{R})$, which is the same as the minimum of $\phi_{\restriction \ell}(A)$ for $\phi_{\restriction \ell} \in \mathcal{P}_{\mathcal{F}_\ell}$. The duality of semidefinite programing (see, e.g., [3, Theorem 4.1.1]) implies that the dual of (8) is the following semidefinite program:

$$\underset{z \in \mathcal{S}_{\mathcal{F}_\ell}}{\text{minimize}} \quad z(A). \tag{9}$$

Therefore, if the semidefinite method finds a proof that for every $\phi \in \mathrm{Hom}^+(\mathcal{A}, \mathbb{R})$ it holds that $\phi(A) \geq c$, where $c \in \mathbb{R}$, it also holds that $z(A) \geq c$ for any $z \in \mathcal{S}_{\mathcal{F}_\ell}$. In particular, if $\phi, \psi \in \mathrm{Hom}^+(\mathcal{A}, \mathbb{R})$ and $\lambda \in (0, 1)$, then

$$\lambda \cdot \phi_{\restriction \ell}(A) + (1 - \lambda) \cdot \psi_{\restriction \ell}(A) \geq c. \tag{10}$$

Note that the $|\mathcal{F}_\ell|$-dimensional vector $\lambda \cdot \phi_{\restriction \ell} + (1 - \lambda) \cdot \psi_{\restriction \ell}$ is usually not a local density ℓ-profile of any convergent sequence of graphs.

References

[1] John-Adrian Bondy and U. S. R. Murty. *Graph theory.* Graduate texts in mathematics. Springer, 2007.

[2] Adolph W. Goodman. On sets of acquaintances and strangers at any party. *Amer. Math. Monthly*, 66(9):778–783, 1959.

[3] Bernd Gärtner and Jiří Matoušek. *Approximation Algorithms and Semidefinite Programming.* Springer, 2012.

[4] Alexander A. Razborov. Flag algebras. *J. Symbolic Logic*, 72(4):1239–1282, 2007.

[5] Andrey Tychonoff. Über die topologische erweiterung von räumen. *Math. Ann.*, 102(1):544–561, 1930.