ssdnm
środowiskowe
studia doktoranckie
z nauk matematycznych

Andrzej Grzesik

Uniwersytet Jagielloński

Flag algebras and Turán problems

Praca semestralna nr 3

(semestr letni 2010/11)

Opiekun pracy: Jarosław Grytczuk

KAPITAŁ LUDZKI
NARODOWA STRATEGIA SPÓJNOŚCI

UNIA EUROPEJSKA
EUROPEJSKI
FUNDUSZ SPOŁECZNY

FLAG ALGEBRAS AND TURÁN PROBLEMS

ANDRZEJ GRZESIK

Abstract. One of the earliest results in extremal combinatorics is Mantel's theorem from 1907 stating that a triangle-free graph on n vertices has at most $n^2/4$ edges. However, similar question posed by Turán in 1941 remains open: what is the maximum number of edges in a 3-uniform hypergraph with no tetrahedron? This question can be considered in wider sense. Given a family of r-uniform hypergraphs $\mathcal{F}$, we want to determine the maximum number of edges in an r-uniform hypergraph on n vertices that does not contain a copy of any hypergraph from family $\mathcal{F}$. Despite substantial efforts, to date there are very few results on this problem, even asymptotically. However, recently Razborov developed an approach to this problem via flag algebras that drastically improves our ability to solve or get approximative results to Turán problems. This article presents intuitions staying behind flag algebras and surveys methods related to this topic.

1. Introduction

An *r-graph* is a pair $G = (V(G), E(G))$, where $V(G)$ is a set of *vertices* and $E(G)$ is a family of r-element subsets of $V(G)$ called *edges*. So, a 2-graph is a simple graph.

We define the *edge density* of an r-graph G as

$$d(G) = \frac{|E(G)|}{\binom{n}{r}}.$$

Given a family $\mathcal{F}$ of r-graphs we say that an r-graph G is *$\mathcal{F}$-free* if G does not contain a subgraph isomorphic to any member of $\mathcal{F}$.

For any integer $n \geq 1$ we define the *Turán number* of $\mathcal{F}$ to be

$$\text{ex}(n, \mathcal{F}) = \max\{|E(G)| : G \text{ is } \mathcal{F}\text{-free}, |V(G)| = n\}.$$

The *Turán density* of $\mathcal{F}$ is defined to be the following limit (it always exists)

$$\pi(\mathcal{F}) = \lim_{n\to\infty} \frac{\text{ex}(n, \mathcal{F})}{\binom{n}{r}}.$$

We can generalize these definitions. Let A be a given r-graph on k vertices ($k \geq r$). We define a set $C_A(G)$ consisting of all k-element subsets of $V(G)$ inducing a subgraph isomorphic to A. Now, we can define a generalized *density* of an r-graph G as

$$d_A(G) = \frac{|C_A(G)|}{\binom{n}{k}},$$

which is just a probability that a random $|V(A)|$-element set of vertices from G induces a graph isomorphic to A. If A is just a single edge, we get definition of typical edge density. When $r = 2$ and A is a triangle, we get so-called triangle density.

We generalize also the definition of the *Turán number* of $\mathcal{F}$ to be

$$\mathrm{ex}_A(n, \mathcal{F}) = \max\{|C_A(G)| : G \text{ is } \mathcal{F}\text{-free, } |V(G)| = n\}.$$

Then, the *Turán density* of $\mathcal{F}$ becomes

$$\pi_A(\mathcal{F}) = \lim_{n\to\infty} \frac{\mathrm{ex}_A(n, \mathcal{F})}{\binom{n}{k}}.$$

This limit always exists because the sequence in the definition forms a decreasing sequence of real numbers in $[0, 1]$.

Determining the Turán density is equivalent to obtaining an asymptotic result $\mathrm{ex}_A(n, \mathcal{F}) \approx \pi_A(\mathcal{F})\binom{n}{k}$, provided that we are in the 'non-degenerate' case when $\pi_A(\mathcal{F}) > 0$.

2. Intuitions

Flag algebras give us systematic approach to counting arguments. They will be fully defined in the next sections, here we present some intuitions. We will focus on Turán density connected with some fixed graph A. Let $\mathcal{F}$ be a family of forbidden r-graphs whose Turán density we wish to compute (or at least approximate).

Let us fix some really large graph G and some small $l \geq r$. Instead of counting appearances of A in graph G, we can count them in each possible graph H on l vertices ($l \geq |V(A)|$), and then, count the number of appearances of H in G. Thus we can write

$$d_A(G) = \sum_{|V(H)|=l} d_H(G) d_A(H).$$

Now, let let $\mathcal{F}_l^0$ be a family of all $\mathcal{F}$-free r-graphs on l vertices up to isomorphism (the reason why we place 0 here will turn out later). If G is $\mathcal{F}$-free, than $d_H(G) = 0$ for H not in $\mathcal{F}_l^0$, so

$$d_A(G) = \sum_{H\in\mathcal{F}_l^0} d_H(G) d_A(H). \tag{1}$$

In particular we have an upper bound

$$d_A(G) \leq \max_{H\in\mathcal{F}_l^0} d_A(H).$$

If l is sufficiently small we can explicitly determine $d_A(H)$ for all $H \in \mathcal{F}_l^0$ by computer search. But this bound on $d_A(G)$ is in general rather poor.

The idea of the method is to generate further inequalities on the probabilities $d_H(G)$ that improve this bound. If we have a linear inequality

$$\sum_{H\in\mathcal{F}_l^0} c_H d_H(G) \geq 0,$$

then

$$d_A(G) \leq \sum_{H\in\mathcal{F}_l^0} d_H(G)(d_A(H)+c_H) \leq \max_{H\in\mathcal{F}_l^0}(d_A(H)+c_H),$$

which may be an improvement if some coefficients c_H are negative.

Moreover, the inequalities we need can be of the form

$$\sum_{H\in\mathcal{F}_l^0} c_H d_H(G) + o(1) \geq 0, \tag{2}$$

where $o(1)$ is taken with respect to $|V(G)|$. By using such inequality we get

$$d_A(G) \leq \max_{H\in\mathcal{F}_l^0}(d_A(H)+c_H) + o(1).$$

Thus

$$\pi_A(\mathcal{F}) \leq \max_{H\in\mathcal{F}_l^0}(d_A(H)+c_H).$$

Let us consider an easy example. We will prove Mantel's theorem stating that triangle-free graph on n vertices has at most $n^2/4$ edges. It is easy to see that it is enough to show that $d(G) \leq 1/2 + o(1)$ for every triangle-free graph G, where $o(1)$ is taken with respect to $n = |V(G)|$.

We will consider $l = 3$. There are only 3 triangle-free graphs on 3 vertices – $\overline{K_3}$, $K_2\cup K_1$ and P_3. For every G on at least 3 vertices, from (1) we get

$$d(G) = d(\overline{K_3})d_{\overline{K_3}}(G) + d(K_2\cup K_1)d_{K_2\cup K_1}(G) + d(P_3)d_{P_3}(G) = \frac{1}{3}d_{K_2\cup K_1}(G) + \frac{2}{3}d_{P_3}(G).$$

If we manage to prove inequality

$$\frac{1}{2}d_{\overline{K_3}}(G) - \frac{1}{6}d_{K_2\cup K_1}(G) - \frac{1}{6}d_{P_3}(G) + o(1) \geq 0, \tag{3}$$

then adding it to the previous equation, we will get

$$d(G) \leq \frac{1}{2}d_{\overline{K_3}}(G) + \frac{1}{6}d_{K_2\cup K_1}(G) + \frac{1}{2}d_{P_3}(G) + o(1) \leq \frac{1}{2} + o(1).$$

This is exactly the inequality we would like to prove.

So, the only thing we need when we are using this technique, is to know how to get inequalities like (3).

Let us focus on one particular vertex in graph G (on pictures we will denote it by unfilled circle). We can define density of a graph with one vertex fixed in similar way as before – as probability of finding an induced copy of this graph with one vertex already fixed. Let G' be a graph G with some vertex fixed. For example, from this definition density $d_{\circ-\bullet}(G')$ is equal to number of vertices connected to fixed vertex divided by $n-1$.

We can write the inequality

$$\left(d_{\circ\,\bullet}(G') - d_{\circ-\bullet}(G')\right)^2 \geq 0$$

and so

$$d_{\circ\,\bullet}(G')d_{\circ\,\bullet}(G') - 2d_{\circ\,\bullet}(G')d_{\circ-\bullet}(G') + d_{\circ-\bullet}(G')d_{\circ-\bullet}(G') \geq 0. \tag{4}$$

The product $d_{\circ\,\bullet}(G')d_{\circ\,\bullet}(G')$ is a probability that two verices chosen at random (we may also chose two times one vertex) are not connected to the fixed vertex. On the other hand, the sum $d_{\bullet\,\circ\,\bullet}(G') + d_{\bullet-\bullet\,\circ}(G')$ represents the same probability, but we assume that we will not chose the same vertex two times. Probability of this event is going to 0 when the size of graph G increases. Hence, we can write

$$d_{\circ\,\bullet}(G')d_{\circ\,\bullet}(G') = d_{\bullet\,\circ\,\bullet}(G') + d_{\bullet-\bullet\,\circ}(G') + o(1).$$

In a similar way, we can prove

$$d_{\circ\,\bullet}(G')d_{\circ-\bullet}(G') = \frac{1}{2}\left(d_{\bullet\ \circ-\bullet}(G') + d_{\bullet-\bullet-\circ}(G')\right) + o(1),$$

$$d_{\circ-\bullet}(G')d_{\circ-\bullet}(G') = d_{\bullet-\circ-\bullet}(G') + o(1).$$

Using these relations, we get

$$d_{\bullet\,\circ\,\bullet}(G') + d_{\bullet-\bullet\,\circ}(G') - d_{\bullet\ \circ-\bullet}(G') - d_{\bullet-\bullet-\circ}(G') + d_{\bullet-\circ-\bullet}(G') + o(1) \geq 0.$$

Now, we can average over all possible choices of the fixed vertex. For example, averaging $d_{\bullet-\bullet\,\circ}(G')$ we will asymptotically get $\frac{1}{3}d_{\bullet-\bullet\,\bullet}(G)$, because only in one case out of three possibilities of chosing a vertex in graph $\bullet-\bullet\,\bullet$ we will get graph $\bullet-\bullet\,\circ$. Thus, we can write

$$d_{\bullet\,\bullet\,\bullet}(G) + \frac{1}{3}d_{\bullet-\bullet\,\bullet}(G) - \frac{2}{3}d_{\bullet-\bullet\,\bullet}(G) - \frac{2}{3}d_{\bullet-\bullet-\bullet}(G) + \frac{1}{3}d_{\bullet-\bullet-\bullet}(G) + o(1) \geq 0$$

and so

$$d_{\bullet\,\bullet\,\bullet}(G) - \frac{1}{3}d_{\bullet-\bullet\,\bullet}(G) - \frac{1}{3}d_{\bullet-\bullet-\bullet}(G) + o(1) \geq 0.$$

Dividing this equality by 2, we get (3), which we wanted to prove.

Summarizing, we started with some non-negative quadratic inequality on densities with some vertices fixed (in this example – inequality (4) with one vertex fixed). Then we changed multiplication of densities into densities of bigger graphs, averaged over all possible choices of fixed vertices, and we obtained inequality of the form (2). The question is, how to get the starting non-negative quadratic inequality containing unknown coefficients and multiplications of densities. Such inequality can be considered as non-negativity of the product of unknown matrix of coefficients by vector of densities (from both sides). If we assume that the unknown matrix is positive semidefinite we will get non-negativity. Thus, we can consider the semidefinite programming problem – minimization of the upper bound with condition that matrix of variables is positive semidefinite. Of course, we can use more than one such inequality.

Flag algebras gives us language to quickly do manipulations (like multiplication or averaging) on densities, like we did in the above example. The idea of the method is clear – we assume non-negativity of some inequalities on densities with some unknown variables and then we make semidefinite programming and find the best coefficients. This inequalities can be of any forms – for example quadratic one, like in the above example (this will be fully described in section 3), taken from Cauchy-Schwarz inequality, or from differentiating. More examples will be presented in section 4.

3. Semidefinite method

In this section we will present one of the systematic approaches to Turán problems.

We define a *type* σ to be an $\mathcal{F}$-free r-graph on s vertices ($s \geq 0$) together with a bijective labelling function $\theta : [s] = \{1, 2, \ldots, s\} \longrightarrow V(\sigma)$. Then we define a *$\sigma$-flag* F to be an $\mathcal{F}$-free r-graph containing an induced copy of σ labelled by θ. We define the order of the flag $|F|$ to be $|V(F)|$.

In other words, if we have a given family $\mathcal{F}$ and a type σ (a graph with all vertices labelled by consecutive numbers from 1 to s) a σ-flag of order m is just an $\mathcal{F}$-free graph on m vertices, which has s labelled vertices inducing σ.

Given a graph G, let us fix a type σ on s vertices, and integers $l > s$, and $m \leq (l+s)/2$. This bound on m ensures that an r-graph on l vertices can contain two subgraphs on m vertices overlapping in exactly s vertices. Let $\mathcal{F}^{\sigma}_{m}$ be the set of all σ-flags of order m, up to isomorphism. Let Θ be the set of all injections from $[s]$ to $V(G)$. Given $F \in \mathcal{F}^{\sigma}_{m}$ and $\theta \in \Theta$, we define induced density of a flag $d_{F,\theta}(G)$ to be the probability that an m-element set V' chosen uniformly at random from $V(G)$ with $\mathrm{im}(\theta) \subseteq V'$, induces a σ-flag that is isomorphic to F. When $s = 0$, that is, if θ is theempty mapping and F is a usual r-graph, definition of $d_{F,\theta}(G)$ coincides with original definition of density $d_F(G)$.

If $F_a, F_b \in \mathcal{F}^{\sigma}_{m}$ and $\theta \in \Theta$, we can define $d_{F_a,F_b}(G;\theta)$ to be the probability that if we choose a random m-element set $V_a \subset V(G)$ with $\mathrm{im}(\theta) \subset V_a$ (so we are choosing only $m - s$ elements), and then we choose a random m-element set $V_b \subset V(G)$ such that $\mathrm{im}(\theta) = V_a \cap V_b$, then induced σ-flags are isomorphic to F_a and F_b, respectively. There is a difference between $d_{F_a,F_b}(G;\theta)$ and the product $d_{F_a}(G;\theta)d_{F_b}(G;\theta)$, because we assume that we cannot choose the same vertices during the second choice. But, when G is large, this probability is negligible, as the following lemma tells us.

Lemma 1 (Razborov [12]). *For any F_a, $F_b \in \mathcal{F}^{\sigma}_{m}$ and $\theta \in \Theta$,*

$$d_{F_a}(G;\theta)d_{F_b}(G;\theta) = d_{F_a,F_b}(G;\theta) + o(1),$$

where the o(1) term tends to 0 as $|V(G)|$ tends to infinity.

Let us consider σ-flags $F_i \in \mathcal{F}_m^\sigma$. Assign some real coefficients a_i to these flags, and for fixed $\theta : [s] \longrightarrow V(G)$ consider the inequality

$$\left(\sum_{F_i \in \mathcal{F}_m^\sigma} a_i d_{F_i}(G;\theta) \right)^2 \geq 0.$$

Expanding the square we have

$$\sum_{F_i, F_j \in \mathcal{F}_m^\sigma} a_i a_j d_{F_i}(G;\theta) d_{F_j}(G;\theta) \geq 0.$$

We can also consider coefficients q_{ij} instead of products $a_i a_j$. If we assume that the matrix $Q = (q_{ij})_{F_i, F_j \in \mathcal{F}_m^\sigma}$ is positive semidefinite, then we will get the inequality

$$\sum_{F_i, F_j \in \mathcal{F}_m^\sigma} q_{ij} d_{F_i}(G;\theta) d_{F_j}(G;\theta) \geq 0.$$

Using the above lemma we obtain

$$\sum_{F_i, F_j \in \mathcal{F}_m^\sigma} q_{ij} d_{F_i, F_j}(G;\theta) + o(1) \geq 0.$$

Now, we average over a uniformly random choice of $\theta \in \Theta$, and using the linearity of expectation we get

$$\sum_{F_i, F_j \in \mathcal{F}_m^\sigma} q_{ij} \mathbb{E}_{\theta \in \Theta}[d_{F_i, F_j}(G;\theta)] + o(1) \geq 0.$$

This expectation can be computed by averaging over all l-vertex subgraphs of G. Let us denote Θ_H as the set of all injective mappings $\theta : [s] \longrightarrow V(H)$ and recall that $\mathcal{F}_l^0$ is the family of all $\mathcal{F}$-free r-graphs on l vertices, up to isomorphism. Thus, we have

$$\mathbb{E}_{\theta \in \Theta}[d_{F_a, F_b}(G;\theta)] = \sum_{H \in \mathcal{F}_l^0} \mathbb{E}_{\theta \in \Theta_H}[d_{F_a, F_b}(H;\theta)] d_H(G).$$

Hence

$$\sum_{H \in \mathcal{F}_l^0} \sum_{F_i, F_j \in \mathcal{F}_m^\sigma} q_{ij} \mathbb{E}_{\theta \in \Theta_H}[d_{F_a, F_b}(H;\theta)] d_H(G) + o(1) \geq 0.$$

Defining $c_H(\sigma, m, Q) = \sum_{F_i, F_j \in \mathcal{F}_m^\sigma} q_{ij} \mathbb{E}_{\theta \in \Theta_H}[d_{F_a, F_b}(H;\theta)]$, we get

$$\sum_{H \in \mathcal{F}_l^0} c_H(\sigma, m, Q) d_H(G) + o(1) \geq 0,$$

which is exactly inequality (2), which can be used to get better upper bound for Turán density.

Furthermore, we can consider t choices of (σ_i, m_i, Q_i), where each σ_i is a type, each $m_i \leq (l + |\sigma_i|)/2$ is an integer, and each Q_i is a positive semidefinite matrix of dimension $|\mathcal{F}^{\sigma_i}_{m_i}| \times |\mathcal{F}^{\sigma_i}_{m_i}|$. For $H \in \mathcal{F}^0_l$ define

$$c_H = \sum_{i=1}^{t} c_H(\sigma_i, m_i, Q_i).$$

Since each Q_i is positive semidefinite matrix, we will get

$$\sum_{H \in \mathcal{F}^0_l} c_H d_H(G) + o(1) \geq 0.$$

If we combine this with

$$d_A(G) = \sum_{H \in \mathcal{F}^0_l} d_A(H) d_H(G),$$

then we will get

$$d_A(G) \leq \sum_{H \in \mathcal{F}^0_l} (d_A(H) + c_H) d_H(G) + o(1).$$

Hence

$$\pi_A(\mathcal{F}) \leq \max_{H \in \mathcal{F}^0_l} (d_A(H) + c_H). \tag{5}$$

Since c_H may be negative, for an appropriate choice of the (σ_i, m_i, Q_i) this bound may be significantly better than the averaging bound given before.

Note that we now have a semidefinite programming problem: given any particular choice of the (σ_i, m_i) find positive semidefinite matrices Q_i which minimize the bound for $\pi_A(\mathcal{F})$ given by (5).

Summarizing the method:
If we want to bound $\pi_A(\mathcal{F})$, we should:

- pick some (not very large) l;
- determine $\mathcal{F}^0_l$ – a set of all $\mathcal{F}$-free r-graphs on l vertices;
- for each $H \in \mathcal{F}^0_l$ compute $d_A(H)$;
- pick some σ (an r-graph with labelled vertices);
- pick some integer $m \leq (l + |\sigma|)/2$;
- determine $\mathcal{F}^\sigma_m$ – a set all $\mathcal{F}$-free σ-flags F_i (r-graphs containing labelled σ, other vertices are not labelled) on m vertices;
- compute $\mathbb{E}_{\theta \in \Theta_H}[d_{F_i,F_j}(H; \theta)]$ for each $H \in \mathcal{F}^0_l$;
- determine functions $c_H = \sum_{F_i, F_j \in \mathcal{F}^\sigma_m} q_{ij} \mathbb{E}_{\theta \in \Theta_H}[d_{F_i,F_j}(H; \theta)]$ over variables q_{ij} forming a matrix Q of dimension $|\mathcal{F}^\sigma_m| \times |\mathcal{F}^\sigma_m|$;
- minimize $\max_{H \in \mathcal{F}^0_l}(d_A(H) + c_H)$ assuming that Q is a positive semidefinite matrix.

This minimum is the bound for $\pi_A(\mathcal{F})$.

To get better bound we can take many triples (σ_i, m_i, Q_i), and for the functions c_H, take the sums of the functions $c_H(\sigma_i, m_i, Q_i)$.

As an example, we can do once more Mantel's theorem. We take $l = 3$. There are three triangle-free graphs on 3 vertices – ,  and . Edge densities of them are equal to 0, 1/3 and 2/3, respectively. We will take type σ consisting of one vertex labelled by 1 and $m = 2$. There are two σ-flags of size 2 –  and  (unfilled vertex is labelled by 1). Required expectations are in the following table

			
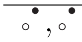	1	1/3	0
	0	2/3	2/3
	0	0	1/3

For an example of this calculation we will calculate the middle column. In graph  we can place label 1 in three possible ways. There are only 2 vertices left, so when we are choosing at random two different vertices, we will just pick these two. And so, in one of the possibilities of placing label 1, we will have pair of two isomorphic flags . In the remaining two possibilities, we will get pair  and . There is no possibility to get pair of two graphs , so there is 0 in the bottom row.

Now, we take variables q_{11}, q_{12}, q_{22} forming symmetric matrix Q. We need to minimize the expression

$$\max\left(q_{11}, \frac{1}{3} + \frac{1}{3}q_{11} + \frac{2}{3}q_{12}, \frac{2}{3} + \frac{2}{3}q_{12} + \frac{1}{3}q_{22}\right),$$

where Q is positive semidefinite. It can be easily seen that it is minimized when

$$Q = \begin{pmatrix} 1/2 & -1/2 \\ -1/2 & 1/2 \end{pmatrix}$$

and the minimum is exactly $1/2$, what we wanted to prove.

The same reasoning was used in [1] for some hypergraph Turán problem, and in [2] to prove that the maximum number of C_5's in a triangle-free graph on n vertices is equal to $(n/5)^5$, which was conjectured by Erdős. This method was also used in many other papers, but with some modifications – it will be more explained in the next section.

4. Algebra

In this section, we would like to present formal definitions of flag algebra, operations in the algebra and also present some useful methods using flag algebras.

In the previous section we defined and used only a type and a flag connected with $\mathcal{F}$-free r-graphs for some forbidden family $\mathcal{F}$. Similar definitions can be made in many other cases, for example for directed graphs or for arbitrary hypergraphs. In general, it can be defined in any universal theory in a first-order language without constants or function symbols. For more model theoretical approach see [9]. Here, we will only use this theory for r-graphs only (for simplicity), but it is worth stressing that everything can be generalized.

The aim when applying flag algebras to Turán problems was to generate useful inequalities of the form $\sum_{F_i \in \mathcal{F}_l^\sigma} b_i d_{F_i}(G) + o(1) \geq 0$ valid for every r-graph $G \in \mathcal{F}^\sigma$ (in particular, when $s = 0$, we have inequalities for graphs, not flags). We can consider $\sum b_i d_{F_i}(G)$ in a little bit different way. Let us take these coefficients b_i and consider formal sum $\sum b_i F_i$ in $\mathbb{R}\mathcal{F}^\sigma$, which is the space of all formal finite linear combinations of σ-flags. We can think of any r-graph $G \in \mathcal{F}^\sigma$ as acting on $\mathbb{R}\mathcal{F}^\sigma$ via the mapping $\sum b_i F_i \longrightarrow \sum b_i d_{F_i}(G)$. So, let us identify such a mapping with G. We can think of every $\mathcal{F}$-free graph as an appropriate mapping on $\mathbb{R}\mathcal{F}^\sigma$. Notice that, by equality $d_{\tilde{F}}(G) = \sum_{F \in \mathcal{F}_l^\sigma} d_{\tilde{F}}(F) d_F(G)$, the linear combination $\tilde{F} - \sum_{F \in \mathcal{F}_l^\sigma} d_{\tilde{F}}(F) F$ is mapped to zero by G. So, we should factor it out. Let $\mathcal{K}^\sigma$ be a linear subspace of $\mathbb{R}\mathcal{F}^\sigma$ generated by all elements of the form

$$\tilde{F} - \sum_{F \in \mathcal{F}_l^\sigma} d_{\tilde{F}}(F) F, \tag{6}$$

where $\tilde{F} \in \mathcal{F}_{\tilde{l}}^\sigma$ and $s \leq \tilde{l} \leq l$. Let $\mathcal{A}^\sigma = \mathbb{R}\mathcal{F}^\sigma / \mathcal{K}^\sigma$ be the quotient space. The only thing left to make algebra is to define multiplication operation. For any $F_1 \in \mathcal{F}_{l_1}^\sigma$ and $F_2 \in \mathcal{F}_{l_2}^\sigma$ we choose any $l \geq l_1 + l_2 - s$ and set

$$F_1 \cdot F_2 = \sum_{F \in \mathcal{F}_l^\sigma} d_{F_1, F_2}(F) F$$

and expand it by linearity. It can be proved (see [12]) that it is well-defined in $\mathcal{A}^\sigma$ (not in $\mathbb{R}\mathcal{F}^\sigma$), it doesn't depend on the choice of l and gives the structure of a commutative algebra in $\mathcal{A}^\sigma$. We know that $d_{F_1,F_2}(G) = d_{F_1}(G) d_{F_2}(G) + o(1)$, so when $|V(G)|$ is large the mapping G is 'approximate homomorphism' from $\mathcal{A}^\sigma$ to $\mathbb{R}$.

To get better understanding of these definitions, we will present some examples:

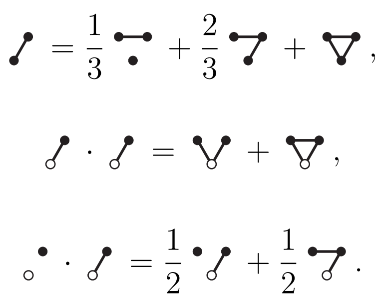

When we are proving something, we are often showing that some inequality holds for particular vertices (for example vertices connected by an edge). And then we are averaging over all possible choices of vertices to get inequality valid for densities, not dependent on any particular vertices. In flag algebras formalism this operation is described as linear operator from $\mathcal{A}^\sigma$ to $\mathcal{A}^0$, which are unlabelled graphs. We define averaging operator of some $F \in \mathcal{F}^\sigma$ to be

$$[\![F]\!] = q_\sigma(F) \cdot F',$$

where F' is an unlabelled version of F, and $q_\sigma(F)$ is the probability that an injective mapping $\theta : [s] \longrightarrow V(F')$ (chosen uniformly at random) defines an induced embedding of σ in F', with the resulting σ-flag isomorphic to F. Next, we extend the operator $[\![\cdot]\!]$ from $\mathcal{F}^\sigma$ to $\mathcal{A}^\sigma$ by linearity.

Some self-explaining examples:

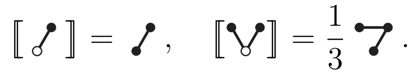

Let D be a path P_3 on 3 edges with first vertex labelled by 1 and adjacent vertex labelled by 2. Then $[\![D]\!] = \frac{1}{6}P_3$, because we can choose an ordered pair of vertices in 12 ways and only 2 of them gives a flag isomorphic to D.

Now, the crucial point of the flag algebras theory goes as follows. With a given σ-flag G we identify the mapping $\mathbb{R}\mathcal{F}^\sigma \longrightarrow \mathbb{R}$ defined before. We can also identify this mapping with vector $(d_F(G))_{F\in\mathcal{F}^\sigma} \in [0,1]^{\mathcal{F}^\sigma}$; and vice versa. The space $[0,1]^{\mathcal{F}^\sigma}$ is compact in the product topology, so any sequence contains a convergent subsequence. Let $\mathrm{Hom}^+(\mathcal{A}^\sigma, \mathbb{R})$ be the set of all homomorphisms ϕ from $\mathcal{A}^\sigma$ to $\mathbb{R}$ such that $\phi(F) \geq 0$ for every $F \in \mathcal{F}^\sigma$. As we noticed, any σ-flag can be identified with such positive homomorphism. It can be proved (see [12]) that for any convergent sequence of σ-flags in $\mathcal{F}^\sigma$ the limit is in $\mathrm{Hom}^+(\mathcal{A}^\sigma, \mathbb{R})$; conversely, any element of $\mathrm{Hom}^+(\mathcal{A}^\sigma, \mathbb{R})$ is the limit of some sequence of σ-flags. This result gives us a correspondence between the final world inequalities $\sum_{F_i\in\mathcal{F}_l^\sigma} b_i d_{F_i}(G) + o(1) \geq 0$ for σ-flags G (in particular unlabelled graphs) and inequalities $\phi(\sum_{F_i\in\mathcal{F}_l^\sigma} b_i F_i) \geq 0$ for $\phi \in \mathrm{Hom}^+(\mathcal{A}^\sigma, \mathbb{R})$. In particular, the Turán density $\pi_A(\mathcal{F}) = \limsup_{G\in\mathcal{F}^0} d_A(G)$ can be rewritten as $\pi_A(\mathcal{F}) = \max_{\phi\in\mathrm{Hom}^+(\mathcal{A}^0,\mathbb{R})} \phi(A)$. From compactness, this maximum is achieved by some extremal homomorphism. In other words, positive homomorphisms are precisely those corresponding to the limits of convergent graph sequences. We will write $F \geq 0$ for σ-flag $F \in \mathcal{A}^\sigma$ if $\phi(F) \geq 0$ for any positive homomorphism ϕ. One can think of the value of $\phi([\![F]\!])$ as the expected value of $\phi(F)$, thus, if $\phi(F) \geq 0$ with probability one, then $\phi([\![F]\!]) \geq 0$.

Let us see how it works on an example. We will prove Mantel's theorem exactly in the same way as in section 2, but in the language of flag algebras. We start from relation

and so

Applying the definition of multiplication, we obtain

The same holds after using averaging operator. It is linear, so

This, after applying the definition of averaging operator, gives

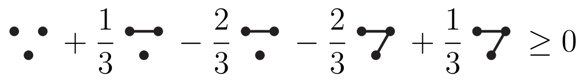

and so

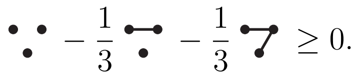

Dividing the last inequality by 2 and adding to equation

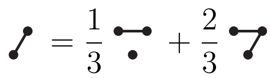

(valid in flag algebra of triangle-free graphs) we get

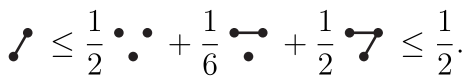

This means that edge density of any triangle-free graph G is at most $1/2 + o(1)$, so $\pi(K_3) \leq 1/2$.

The big 'source' for inequalities on flag algebras are the Cauchy-Schwarz inequalities.

Theorem 2 (Razborov [12]). *For any f, $g \in \mathcal{A}^\sigma$*

$$[\![f^2]\!] \cdot [\![g^2]\!] \geq [\![fg]\!]^2.$$

In particular

$$[\![f^2]\!] \cdot [\![\sigma]\!] \geq [\![f]\!]^2,$$

which implies

$$[\![f^2]\!] \geq 0.$$

As an easy example, we will very quickly show Mantel's theorem. Let us notice that in flag algebra of graphs (without any forbidden graphs) we have

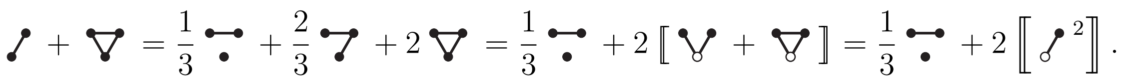

From the theorem , so

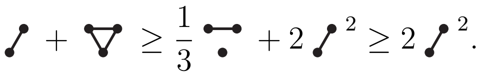

It means that 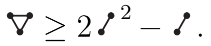. In particular, if edge density is greater than 1/2, then triangle density is greater than 0, so there are some triangles.

We get it more quickly than before, and it is also a stronger result. Of course, it is possible to translate this proof to 'normal' language, and instead of using operations on flags, to make some counting arguments. But, in general, some proofs on flag algebras are using so many inequalities that after translation, it would be unreadable.

Moreover, using flag algebras we can make systematic approach to Turán problems. To bound some $\pi_A(\mathcal{F})$ we can choose some type σ (or few different types), pick some not very large l and make a computer program to generate all possible Cauchy-Schwarz inequalities for flags up to l vertices. Then expand each inequality and, using averaging operator, express them as linear inequalities on graphs on l vertices. After that, formulate a semidefinite programming problem to calculate numerically the best bound. If we are lucky, we can get good bound. After rationalization our proof is numerically stable. This method was used in many papers to obtain new results (see for example [5], [13]).

Flag algebras were also used in problems with different setting than Turán problems. For example to get new results on the Caccetta-Häggkvist Conjecture for triangles (see [3], [9]), where we have assumption about minimal outdegree, or on the problem of selecting heavily covered points ([8]).

It is worth mentioning that flag algebras permit 'differential methods' (see [12]). Maximum value of $\pi_A(\mathcal{F})$ is achieved for some extremal positive homomorphism ϕ, so any small perturbation of ϕ must reduce $\phi(A)$. Perturbation with respect to a single vertex is analogous to some deletion arguments, but general perturbation do not have any obvious analogue in final setting. That is why flag algebras may turn out to be very powerful tool for a wider collections of problems in combinatorics.

References

[1] R. Baber, J. Talbot, *Hypergraphs Do Jump*, Combinatorics, Probability and Computing **20** (2011), 161-171.

[2] A. Grzesik, *On the maximum number of C_5's in a triangle-free graph*, arXiv:1102.0962, 2011.

[3] J. Hladký, D. Král', S. Norine, *Counting flags in triangle-free digraphs*, proceedings of Eurocomb'09, Electronic Notes in Discrete Mathematics **34** (2009), 621-625

[4] H. Hatami, J. Hladký, D. Král', S. Norine, A.A. Razborov, *Non-three-colorable common graphs exist*, arXiv:1105.0307, 2011.

[5] H. Hatami, J. Hladký, D. Král', S. Norine, A.A. Razborov, *On the number of pentagons in triangle-free graphs*, arXiv:1102.1634, 2011.

[6] H. Hatami, S. Norine, *Undecidability of linear inequalities in graph homomorphism densities*, Journal of the American Mathematical Society **24** (2011), 547-565.

[7] P. Keevash, *Hypergraph Turán Problems*, to appear in *Surveys in Combinatorics*, 2011.

[8] D. Král', L. Mach, J.-S. Sereni, *A new lower bound based on Gromov's method of selecting heavily covered points*, arXiv:1108.0297, 2011.

[9] A.A. Razborov, *On the Caccetta-Häggkvist Conjecture with Forbidden Subgraphs*, arXiv:1107.2247, 2011.

[10] A.A. Razborov, *On the Fon-der-Flaass interpretation of extremal examples for Turán's (3,4)-problem*, arXiv:1008.4707, 2010.

[11] A.A. Razborov, *On 3-hypergraphs with forbidden 4-vertex configurations*, SIAM Journal on Discrete Mathematics **24** (2010), 946–963.

[12] A.A. Razborov, *Flag Algebras*, Journal of Symbolic Logic **72** (2007), 1239–1282.

[13] K. Sperfeld, *On the minimal monochromatic K_4-density*, arXiv:1106.1030, 2011.

Theoretical Computer Science, Jagiellonian University, Kraków, Poland
E-mail address: Andrzej.Grzesik@uj.edu.pl