Containment restrictions

Tibor Szabó

Extremal Combinatorics, FU Berlin, WiSe 2017–18

In this chapter we switch from studying constraints on the set operation *intersection*, to constraints on the set relation *containment*. While of course the two are strongly related ($A \subseteq B$ if and only if $A \cap B^c = \emptyset$), the flavour of the problems will be quite different.

The origon of all results in this chapter is Sperner's Theorem, which was proved in the Classics section. Just to recall, it stated that the size of the largest antichain in the Boolean poset is equal to size of (one of) the middle layer(s), that is, $\binom{n}{\lfloor n/2 \rfloor}$. Sperner's Theorem is a consequence of a stronger, weighted inequality called the LYM Inequality, stating that

$$\sum_{F \in \mathcal{F}} \frac{1}{\binom{n}{|F|}} \leq 1,$$

for any antichain $\mathcal{F} \subseteq 2^{[n]}$,

In this section first we discuss a couple of beautiful "ground-set-independent" generalizations of Sperner's Inequality, which were motivated by extremal graph theory problems. In the second section we will start off with the characterization of extremal constructions for Sperner's Theorem.

1 Set-pair inequalities

1.1 The Bollobás Set-Pair Inequality

Theorem 1.1 (Bollobás, 1965). *Let $A_1, A_2, \ldots, A_m$ and $B_1, B_2, \ldots, B_m$ be two sequences of finite sets, such that:*

(i) For all $1 \leq i \leq m$, $A_i \cap B_i = \emptyset$.

(ii) For any $i \neq j$, $A_i \cap B_j \neq \emptyset$.

Then

$$\sum_{i=1}^{m} \frac{1}{\binom{|A_i|+|B_i|}{|A_i|}} \leq 1.$$

Corollary 1.2 (Uniform version). *If additionally to (i) and (ii), we have*

(iii) $|A_i| = k$ and $|B_i| = \ell$,

then $m \leq \binom{k+\ell}{k}$.

Claim 1.3. *Theorem 1.1 implies the LYM inequality.*

Proof. Let $\mathcal{F} = \{F_1, \ldots, F_m\} \subseteq 2^{[n]}$ be an antichain. Set $A_i = F_i$ and $B_i = [n] \setminus F_i = F_i^c$. To verify conditions (i) and (ii) note that $A_i \cap B_j = F_i \cap F_j^c = \emptyset$ is equivalent to $F_i \subseteq F_j$, which happens if and only if $i = j$ since $\mathcal{F}$ is an antichain. Therefore, by Theorem 1.1, we have

$$\sum_{F \in \mathcal{F}} \frac{1}{\binom{n}{|F|}} = \sum_{i=1}^{m} \frac{1}{\binom{|A_i|+|B_i|}{|A_i|}} \leq 1,$$

where the equality follows from the fact that $|F| + |F^c| = n$. $\square$

Our proof of the Bollobás set-pair inequality is a generalization of the double-counting proof we gave for Sperner's Theorem in Chapter **??**. Not to get bored, here we formulate our argument in probabilistic terms.

Proof of Theorem 1.1. Let $X = \bigcup_{i=1}^{m} (A_i \cup B_i)$ be our (finite) ground set. Let $\pi : X \to [|X|]$ be a uniformly random permutation of X. Let $\mathcal{E}_i$ be the event

$$\mathcal{E}_i = \{\pi : \max \pi(A_i) < \min \pi(B_i)\},$$

that is, that the last element of A_i appears before the first element of B_i. Clearly, the event $\mathcal{E}_i$ only depends only on the relative positions of $\pi(A_i)$ and $\pi(B_i)$ within $\pi(A_i \cup B_i)$ and has nothing to with any other value of π (in fact, the ground-set-freeness of the theorem is made possible by this). Formally, conditioning on *any fixed* injection $\pi' : X \setminus (A_i \cup B_i) \to [|X|]$ providing the values of π on $X \setminus (A_i \cup B_i)$, the probability of $\mathcal{E}_i$ is always the same:

$$\mathbb{P}\left(\mathcal{E}_i \mid \pi|_{X \setminus (A_i \cup B_i)} = \pi'\right) = \frac{|A_i|! \cdot |B_i|!}{(|A_i| + |B_i|)!}.$$

This is because $(|A_i| + |B_i|)!$ is the number of ways π' can be extended to a permutation of X, and the number of ways the elements of A_i (or B_i) can be permuted in the first $|A_i|$ (or last $|B_i|$) positions of $[|X|] \setminus \pi'(X \setminus (A_i \cup B_i))$ is equal to $|A_i|!$ (or $|B_i|!$). Hence $\mathbb{P}(\mathcal{E}_i)$ is also equal to $\frac{|A_i|!|B_i|!}{(|A_i|+|B_i|)!}$.

We show that the events $\mathcal{E}_i, 1 \leq i \leq m$, are pairwise disjoint. Let $i \neq j$. Since by (ii) there is an element $x \in A_i \cap B_j$, we have that $\min \pi(B_i) \leq \pi(x) \leq \max \pi(A_j)$. Similarly $\min \pi(B_j) \leq \max \pi(A_i)$. Consequently, if $\mathcal{E}_i$ happens then

$$\min \pi(B_j) \leq \max \pi(A_i) < \min \pi(B_i) \leq \max \pi(A_j),$$

so $\mathcal{E}_j$ does not happen.

Then the probability of the union is the sum of the individual probabilities:

$$1 \geq \mathbb{P}\left(\bigcup_{i=1}^{m} \mathcal{E}_i\right) = \sum_{i=1}^{m} \mathbb{P}(\mathcal{E}_i) = \sum_{i=1}^{m} \frac{|A_i|!\,|B_i|!}{(|A_i| + |B_i|)!} = \sum_{i=1}^{m} \frac{1}{\binom{|A_i|+|B_i|}{|A_i|}}.$$

$\square$

1.1.1 An application: saturated graphs

The Bollobás set pair inequality was motivated by an extremal problem arising from the algorithmic task of constructing an H-free graph on n vertices, with as many edges possible. Given an integer n, a "forbidden" k-uniform hypergraph $\mathcal{H}$, and an ordering of the edges of the complete k-graph $\binom{[n]}{k}$, the Greedy Algorithm goes thorugh the k-sets in the given order, always adding the next k-set e to the k-graph $\mathcal{G}$ it maintains, if this addition does not create a copy of $\mathcal{H}$. This edge-addition rule ensures that the output $\mathcal{G}$ of the Greedy Algorithm is an $\mathcal{H}$-free k-graph on n vertices. The question is *how many* k-edges the output has.

If the input ordering happens to start with the the edges of an $\mathcal{H}$-free graph with the maximum number $ex(n,\mathcal{H})$ of edges, then the algorithm will of course output exactly this k-graph. What can we say about how *bad* the Greedy Algorithm can do? How *small* number of edges can an output graph have?

A k-graph $\mathcal{G}$ is called *$\mathcal{H}$-saturated* if it is $\mathcal{H}$-free, but $\mathcal{G} \cup \{e\}$ contains a copy of $\mathcal{H}$ for every k-set $e \in \binom{[n]}{k} \setminus \mathcal{G}$. Note that a k-graph is $\mathcal{H}$-saturated if and only if it is the output of the Greedy Algorithm for some input ordering of $\binom{[n]}{k}$.

To measure how bad the Greedy Algorithm can do, we introduce the *saturation number* of the k-graph $\mathcal{H}$ as follows,

$$sat(n,\mathcal{H}) := \min\{|\mathcal{G}| : \mathcal{G} \text{ is } \mathcal{H}\text{-saturated } k\text{-graph on } n \text{ vertices}\}.$$

Note that any $\mathcal{H}$-free k-graph with $ex(n,\mathcal{H})$ edges is necessarily $\mathcal{H}$-saturated, as the addition of any non-edge makes the number of edges larger than $ex(n,\mathcal{H})$, so by the definition of the Turán number there must be a copy of $\mathcal{H}$. Hence $sat(n,\mathcal{H}) \leq ex(n,\mathcal{H})$, as sat is the minimum while ex is the maximum number of edges over the same family of graphs.

Example: For the 2-graph K_3 the Turán number is $\lfloor n^2/4 \rfloor$, that is quadratic in n. How small is the satuaration number? The star graph $K_{1,n-1}$ is K_3-saturated and shows that $sat(n,K_3)$ is not more than linear. In fact it is equal to $n-1$, since any graph with $n-2$ is disconnected, so an edge connecting two components can be added and such an edge does not create a copy of K_3.

The star construction for K_3 can be generalized to arbitary k-uniform t-cliques. Considering a set $T \subseteq [n]$ of $t-k$ vertices, form the k-graph $\mathcal{G} := \{K : |K| = k, |K \cap T| \neq \emptyset\}$ of all k-sets having a non-empty intersection with T. Such a k-graph is clearly $K_t^{(k)}$-free, since every t-set has at least k vertices outside of T and these form a k-set that is not in $\mathcal{G}$. For every $e \notin \mathcal{G}$ however, $T \cup e$ is a t-set with all its k-subsets, except e, contained in $\mathcal{G}$. So $\mathcal{G}$ is $K_t^{(k)}$-saturated. This generalized star construction contains $\binom{n}{k} - \binom{n-t+k}{k}$ hyperedges, since exactly those k-sets are not included in it that are contained in the complement of T.

Next we use the set-pair inequality to show that this upper bound on $sat(n,K_t^{(k)})$ is the best possible for all values of the parameters.

Theorem 1.4 (Bollobás, 1965)**.** *For every $n \geq t \geq k \geq 2$, we have that*

$$sat(n,K_t^{(k)}) = \binom{n}{k} - \binom{n-t+k}{k}.$$

Proof. For the $\geq$-direction, let $\mathcal{G}$ be a $K_t^{(k)}$-saturated k-graph on n vertices. Let $e_1,\ldots,e_m \in \binom{[n]}{k} \setminus \mathcal{G}$ be a list of all k-sets *not* in $\mathcal{G}$. Then the addition any e_i to $\mathcal{G}$ creates a copy of $K_t^{(k)}$; let K_i be the vertex set of one of these t-cliques.

We use Corollary 1.2 with $A_i = e_i$ and $B_i = [n] \setminus K_i$. The condition $A_i \cap B_i = \emptyset$ holds, since $e_i \subseteq K_i$. For $i \neq j$ we have $A_i \cap B_j = e_i \cap ([n] \setminus K_j) \neq \emptyset$. Indeed, otherwise $K_j \supseteq e_i$ in which case $\mathcal{G}$ restricted to the vertex set K_j is missing not only the k-set e_j, but also e_i, a contradiction.

Since $|e_i| = k$ and $|K_i| = n-t$ for every $i =, 1\ldots,m$, Corollary 1.2 is applicable and gives $m \leq \binom{k+n-t}{k}$. □

A surprising feature of the Theorem is the precise determination of the minimum number of edges in an n-vertex $K_t^{(k)}$-saturated graphs for *every value* of the parameters. Compare this with our knowledge when we aim to determine the *maximum* instead: the Turán number is not know, even asymptotically, for any $n \geq t \geq k \geq 3$.

For $k = 2$, i.e. for the graph case, the Turán number is known precisely and its value is *quadratic* in n for every constant $t \geq 3$. The saturataion number on the other hand is $\binom{n}{2} - \binom{n-t+2}{2} = ...$ is only linear in n. So the Greedy Algorithm can perform *very badly* in comparison to the maximum.

1.2 The Lovász Set-Pair Inequality

1.2.1 Weakly saturated graphs

An alternative greedy way to construct $\mathcal{H}$-free k-graphs is to do everything backwards: start from the complete k-graph on n vertices, go through its k-sets in a given order, leaving out the next edge if it is contained in a copy of $\mathcal{H}$. After this Reverse Greedy Algorithm is finished considering all edges, the resulting k-graph is $\mathcal{H}$-free.

How well does this algorithm perform? Again, it depends on the input ordering of the k-sets. Putting the edges of an $\mathcal{H}$-free graph with $ex(n, \mathcal{H})$ edges to the end of the order ensures that the algorithm outputs exactly this extremal graph. How bad can the Reverse Greedy Algorithm do? A short moment of contemplation will convince us that RGA cannot perform better in the worst case than GA.

A k-graph $\mathcal{G}$ on vertex set $[n]$ is called *weakly $\mathcal{H}$-saturated* if it is $\mathcal{H}$-free and there exists a sequence $e_1, \dots, e_m \in \binom{[n]}{k} \setminus \mathcal{G}$ of the k-sets not in $\mathcal{G}$, such that for every $i = 1, \dots, m$ the addition of e_i to $\mathcal{G} \cup \{e_1, \dots, e_{i-1}\}$ creates a new copy of $\mathcal{H}$.

An $\mathcal{H}$-saturated k-graph $\mathcal{G}$ is also weakly $\mathcal{H}$-saturated, since for *any* ordering $e_1, \dots, e_m$ of the k-sets in $\binom{[n]}{k} \setminus \mathcal{G}$ each k-set e_i creates a copy of $\mathcal{H}$ already together with just the edges in $\mathcal{G}$ (and it does not need the "help" of the further k-sets $e_1, \dots, e_{i-1}$).

Claim 1.5. *A k-graph $\mathcal{G}$ on vertex set $[n]$ is the output of the Reverse Greedy Algorithm for some ordering of the edges if and only if it is weakly $\mathcal{H}$-saturated.*

Proof. If $\mathcal{G}$ is weakly $\mathcal{H}$-saturated, then let $e_1, \dots, e_m$ be the appropriate ordering of the k-sets outside of $\mathcal{G}$, given by the definition. Putting these k-sets at the beginning of the input ordering in *reverse* order will make RGA to delete each of them, since each e_i participates in a copy of $\mathcal{H}$ together with the edges in $\mathcal{G} \cup \{e_1, \dots, e_{i-1}\}$. After deleting all these k-sets, RGA ends up with $\mathcal{G}$ and does not delete any more k-sets, since $\mathcal{G}$ is $\mathcal{H}$-free.

Let now $\mathcal{G}$ be a k-graph that is the output of RGA for some ordering of $\binom{[n]}{k}$. Let $e_1, \dots, e_m$ be the edges that were deleted by RGA, in *reverse order.* Then the reason e_i was deleted by RGA is that it participated in a copy of $\mathcal{H}$ with the edges still present. These edges are exactly $\mathcal{G} \cup \{e_1, \dots, e_{i-1}\}$, since $e_m, \dots, e_{i+1}$ were already deleted and the edges of $\mathcal{G}$ are kept throughout. Hence $\mathcal{G}$ weakly $\mathcal{H}$-saturated. $\square$

Consequently the worst case behaviour of the Reverse Greedy Algorithm can be measured by the extremal number

$$wsat(n, H) = \min\{|\mathcal{G}| : \mathcal{G} \text{ is weakly } \mathcal{H}\text{-saturated on } n \text{ vertices}\}.$$

By the above

$$wsat(n, \mathcal{H}) \leq sat(n, \mathcal{H}) \leq ex(n, \mathcal{H}).$$

Example. For $H = K_3$ the weak saturation number is the same as the saturation number $n - 1$, because any graph with less than $n - 1$ edges is disconnected. Indeed, if G was disconnected and weakly K_3-saturated, with an appropriate ordering $e_1, \dots e_m$ of its non-edges, then the first edge e_i between two of G's components does not create a K_3 in $E(G) \cup \{e_1, \dots, e_{i-1}\}$, a contradiction.

It turns out that the two saturation numbers of cliques are equal for all values of the parameters.

Theorem 1.6 (Lovász, 197?)**.** *For every $n \geq t \geq k \geq 2$, we have*

$$sat(n, K_t^{(k)}) = wsat(n, K_t^{(k)}).$$

Proof. Let $\mathcal{G}$ be a weakly $K_t^{(k)}$-saturated k-graph on n vertices and let $e_1, \ldots, e_m$ be an ordering of the k-sets in $\binom{[n]}{k} \setminus \mathcal{G}$, such that the addition of e_i to $\mathcal{G} \cup \{e_1, \ldots, e_{i-1}\}$ creates a new copy of $K_t^{(k)}$. Let K_i be the vertex set of this t-clique.

We try to mimic the proof of Theorem 1.4 using set pairs and see if anything goes wrong. We define $A_i = e_i$ and $B_i = [n] \setminus K_i$ for every $i = 1, \ldots, m$. Then again $e_i \subseteq K_i$, so $A_i \cap B_i = \emptyset$. Now, however, $A_i \cap B_j = \emptyset$, that is $e_i \subseteq K_j$, does not immediately lead to a contradiction for every $i \neq j$. For $j > i$ for example, there is nothing wrong with the k-set e_i being part of the new t-clique created by e_j on K_j, since it is using edges of $\mathcal{G} \cup \{e_1, \ldots, e_{j-1}\}$, of which e_i is part of. We can conclude a contradiction though for $j < i$, when e_j should create the new t-clique on K_j together with some edges of $\mathcal{G} \cup \{e_1, \ldots, e_{j-1}\}$, of which e_i is *not* part of.

We resolve this situation with a new skew set-pair criterion, Theorem 1.7, which gives the same conclusion as Corollary 1.2, but under a weaker assumption (ii), requiring that $A_i \cap B_j \neq \emptyset$ for every $j > i$ only. $\square$

Theorem 1.7 (Lovász, 197?). *Let $A_1, A_2, \ldots, A_m$ and $B_1, B_2, \ldots, B_m$ be two sequences of finite sets, such that:*

1. $|A_i| = k$ *and* $|B_i| = \ell$ *for every* $i = 1, \ldots, m$,
2. $A_i \cap B_i = \emptyset$ *for all* $1 \leq i \leq m$,
3. $A_i \cap B_j \neq \emptyset$, *for any* $i > j$,

Then

$$m \leq \binom{k+\ell}{k}.$$

Remark. The uniformity assumption in the statement is essential. The skew analogue of Theorem 1.1 is not true. (HW)

The statement of Theorem 1.7 was conjectured already by Bollobás, but its proof took some time to be found. Most likely since it probably *had to be* vastly different from the combinatorial argument for Theorem 1.1.

1.2.2 Proof of the Lovász set-pair inequality

The main challenge, again, is to accommodate the ground-set-free nature of the statement. In the proof of Theorem 1.1 this was achieved via conditioning in the probability space. Here we give a proof using linear algebra. In our linear algebraic proofs so far we have associated coordinates with the vertices and gave a bound on the dimension that obviously had to depend on the number of coordinates. Now we will work in a space with dimension depending on the uniformity k, and identify the elements of our ground set $\cup_{i=1}^m (A_i \cup B_i) =: X$ with vectors such that any $k+1$ of them are linearly independent.

A set of vectors $S \subseteq \mathbb{F}^d$ is said to be in *general linear position* if any d vectors from S is linearly independent.

Within a set of vectors in general linear position there is no linear dependence beyond what is absolutely necessary due to the dimension of the space these vectors live in. Fortunately for any dimension d there exists a large set of vectors that is in general linear position. In particular the point set $\{(1, \alpha, \alpha^2, \ldots, \alpha^{d-1}) \in \mathbb{F}^d : \alpha \in \mathbb{F}\}$, the so-called *moment curve*, provides a set of $|\mathbb{F}|$ vectors in general linear position. Any d vectors $(1, \alpha_i, \alpha_i^2, \ldots, \alpha_i^{d-1})$, $i = 1, \ldots, d$, from the moment curve are linearly independent, since the determinant of the matrix formed by them as row vectors is $\prod_{i<j}(\alpha_i - \alpha_j) \neq 0$ (the Vandermonde determinant).

Proof of Theorem 1.7. We will work over the field of reals, so it is possible to associate a vector $\mathbf{v_x} \in \mathbb{R}^{k+1}$ to each vertex $x \in X$. For each A_i we consider the subspace V_i generated by the vectors $\mathbf{v_x}$ associated with the elements x of A_i. Due to the general linear position of these vectors, the dimension $\dim V_i = k$. Hence the orthogonal complement $V_i^\perp$ is 1-dimensional. We associate with the index $i \in [m]$ an arbitrary non-zero vector $\mathbf{u_i}$ from this orthogonal complement. It turns out that orthogonality to $\mathbf{u_i}$ characterizes containment in A_i.

Key Observation $\mathbf{u_j} \cdot \mathbf{v_y} = 0$ *if and only if* $y \in A_j$.

Indeed, $\mathbf{u_j} \cdot \mathbf{v_y} = 0 \iff \mathbf{v_y} \in V_j \iff \{\mathbf{v_x} : x \in A_j\} \cup \{\mathbf{v_y}\}$ is linearly dependent $\iff y \in A_j$ (otherwise the linear dependence of the vectors $\mathbf{v_y}$ and $\mathbf{v_x}, x \in A_j$ would contradict the general linear position assumption).

For each i we define the polynomials

$$f_i(\mathbf{x}) = \prod_{y \in B_i} \mathbf{x} \cdot \mathbf{v_y} = \prod_{y \in B_i} \left((\mathbf{v_y})_1 \mathbf{x}_1 + \cdots + (\mathbf{v_y})_{k+1} \mathbf{x}_{k+1} \right).$$

For these polynomials f_i and the vectors $\mathbf{u_i}$, the determinant criterion (Lemma **??**) hold. The substitution $f_i(\mathbf{u_j}) = 0$ if and only if there is a $y \in B_i$ such that $\mathbf{u_j} \cdot \mathbf{v_y} = 0$. This, by our Key Observation, is equivalent with the existence of $y \in B_i$ for which also $y \in A_j$. In other words, if $B_i \cap A_j \neq \emptyset$.

So the skew conditions of our theorem translates to the substitution matrix $(f_i(\mathbf{u_j}))_{i,j}$ being upper triangular with non-zero diagonal entries. Lemma **??** then implies the linear independence of the polynomials $f_1, \ldots, f_m$.

These polynomials live in the space of homogeneous polynomials of degree ℓ. Each such polynomial is the sum of terms of the form $c_{\alpha_1,\ldots,\alpha_{k+1}} \prod_{i=1}^{k+1} x_i^{\alpha_i}$, where $\alpha_1 + \cdots + \alpha_{k+1} = \ell$. This is exactly the problem of partitioning ℓ indistinguishable items into $k+1$ distinguishable boxes. The number of such partitions is $\binom{\ell+k+1-1}{\ell}$. This is the dimension of our space of polynomials, and hence it is an upper bound on the number m of linearly independent polynomials we found in the space. $\square$

Remark Considering that Theorem 1.7 is also a generalization of Sperner's Theorem, we now have proved this fundmental result using combinatorial, probabilistic, and linear algebraic concepts.