11 Containers

Many problems in combinatorics can be phrased in terms of independent sets in hypergraphs.

For example, here is a model question:

Question 11.0.1
How many triangle-free graphs are there on n vertices?

By taking all subgraphs of $K_{n/2,n/2}$, we obtain $2^{n^2/4}$ such graphs. It turns out this gives the correct exponential asymptotic.

Theorem 11.0.2 (Erdős, Kleitman, and Rothschild 1973)
The number of triangle-free graphs on n vertices is $2^{n^2/4+o(n^2)}$.

***Remark* 11.0.3.** It does not matter here whether we consider vertices to be labeled, it affects the answer up to a factor of at most $n! = e^{O(n \log n)}$.

***Remark* 11.0.4.** Actually the original Erdős–Kleitman–Rothschild paper showed an even stronger result: $1 - o(1)$ fraction of all n-vertex triangle-free graphs are bipartite. The above asymptotic can be then easily deduced by counting subgraphs of complete bipartite graphs. The container methods discussed in this section are not strong enough to prove this finer claim.

We can convert this asymptotic enumeration problem into a problem about independent sets in a 3-uniform hypergraph H:

- $V(H) = \binom{[n]}{2}$
- The edges of H are triples of the form $\{xy, xz, yx\}$, i.e., triangles

We then have the correspondence:

- A subset of $V(H)$ = a graph on vertex set $[n]$
- An independent set of $V(H)$ = a triangle-free graph

(Here an ***independent set*** in a hypergraph is a subset of vertices containing no edges.)

Naively applying first moment/union bound does not work—there are too many events to union bound over.

For example, Mantel's theorem tell us the maximum number of edges in an n-vertex triangle-free graph is $\lfloor n^2/4 \rfloor$, obtained by $K_{\lfloor n/2 \rfloor, \lceil n/2 \rceil}$. With a fixed triangle-free graph G, the number of subgraphs of G is $2^{e(G)}$, and each of them is triangle-free. Perhaps we could union bound over all maximal triangle-free graphs? It turns out that there are $2^{n^2/8+o(n^2)}$ such maximal triangle-free graphs, so a union bound would be too wasteful.

In many applications, independent sets are clustered into relatively few highly correlated sets. In the case of triangle-free graphs, each maximal triangle-free graph is very "close" to many other maximal triangle-free graphs.

Is there a more efficient union bound that takes account of the clustering of independent sets?

The container method does exactly that. Given some hypergraph with controlled degrees, one can find a collection of ***containers*** satisfying the following properties:

- Each container is a subset of vertices of the hypergraph.
- Every independent set of the hypergraph is a subset of some container.
- The total number of containers in the collection is relatively small.
- Each container is not too large (in fact, not too much larger than the maximum size of an independent set)

We can then union bound over all such containers. If the number of containers is not too small, then the union bound is not too lossy.

Here are some of the most typical and important applications of the container method:

- Asymptotic enumerations:
  - Counting H-free graphs on n vertices
  - Counting H-free graphs on n vertices and m edges
  - Counting k-AP-free subsets of $[n]$ of size m
- Extremal and Ramsey results in random structures:
  - The maximum number of edges in an H-free subgraph of $G(n, p)$
  - Szemerédi's theorem in a p-random subset of $[n]$
- List coloring in graphs/hypergraphs

The method of hypergraph containers is one of the most exciting developments in this past decade. Some references and further reading:

- The graph container method was developed by Kleitman and Winston (1982) (for counting C_4-free graphs) and Sapozhenko (2001) (for bounding the number of independent sets in a regular graph, giving an earlier version of Theorem 10.4.12)
- The hypergraph container theorem was proved independently by Balogh, Morris, and Samotij (2015), and Saxton and Thomason (2015).
- See the 2018 ICM survey of Balogh, Morris, and Samotij for an introduction to the topic along with many applications
- See Samotij's survey article (2015) for an introduction to the graph container method
- See Morris' lecture notes (2016) for a gentle introduction to the proof and applications of hypergraph containers.

11.1 Containers for triangle-free graphs

The number of triangle-free graphs

We will prove Theorem 11.0.2 that the number of triangle-free graphs on n vertices is $2^{n^2/4+o(n^2)}$.

Theorem 11.1.1 (Containers for triangle-free graphs)
For every $\varepsilon > 0$, there exists $C > 0$ such that the following holds.

For every n, there is a collection $\mathcal{C}$ of graphs on n vertices, with

$$|\mathcal{C}| \le n^{Cn^{3/2}}$$

such that

(a) every $G \in \mathcal{C}$ has at most $(\frac{1}{4} + \varepsilon)n^2$ edges, and

(b) every triangle-free graph is contained in some $G \in \mathcal{C}$.

Proof of upper bound of Theorem 11.0.2. We want to show that the number of n-vertex triangle-free graphs is at most $2^{n^2/4+o(n^2)}$. Let $\varepsilon > 0$ be any real number (arbitrarily small). Let $\mathcal{C}$ be produced by Theorem 11.1.1.

Then every $G \in \mathcal{C}$ has at most $(\frac{1}{4} + \varepsilon)n^2$ edges, and every triangle-free graph is contained in some $G \in \mathcal{C}$. Hence the number of triangle-free graphs is

$$|\mathcal{C}|\, 2^{(\frac{1}{4}+\delta)n^2} \le 2^{(\frac{1}{4}+\varepsilon)n^2 + O_\varepsilon(n^{3/2}\log n)}.$$

Since $\varepsilon > 0$ can be made arbitrarily small, the number triangle-free graphs on n vertices is $2^{(\frac{1}{4}+o(1))n^2}$. ◻

The same proof technique, with an appropriate container theorem, can be used to count H-free graphs.

We write $\operatorname{ex}(n, H)$ for the maximum number of edges in an n-vertex graph without H as a subgraph.

Theorem 11.1.2 (Erdős–Stone–Simonovits)
Fix a non-bipartite graph H. Then

$$\operatorname{ex}(n, H) = \left(1 - \frac{1}{\chi(H) - 1} + o(1)\right)\binom{n}{2}.$$

Note that for bipartite graphs H, the above theorem just says $o(n^2)$, though more precise estimates are available. Although we do not know the asymptotic for $\operatorname{ex}(n, H)$ for all H, e.g., it is still open for $H = K_{4,4}$ and $H = C_8$.

Theorem 11.1.3
Fix a non-bipartite graph H. Then the number of H-free graphs on n vertices is $2^{(1+o(1))\operatorname{ex}(n,H)}$.

The analogous statement for bipartite graphs is false. The following conjecture remains of great interest, and it is known for certain graphs, e.g., $H = C_4$.

Conjecture 11.1.4
Fix a bipartite graph H with a cycle. The number of H-free graphs on n vertices is $2^{O(\operatorname{ex}(n,H))}$.

Mantel's theorem in random graphs

Theorem 11.1.5
If $p \gg 1/\sqrt{n}$, then with probability $1 - o(1)$, every triangle-free subgraph of $G(n, p)$ has at most $(\frac{1}{4} + o(1))pn^2$ edges.

***Remark* 11.1.6.** In fact, a much stronger result is true: the triangle-free subgraph of $G(n, p)$ with the maximum number of edges is whp bipartite (DeMarco and Kahn 2015).

***Remark* 11.1.7.** The statement is false for $p \ll 1/\sqrt{n}$. Indeed, in this case, then the expected number of triangles is $O(n^3 p^3)$, whereas there are whp $n^2 p/2$ edges, and $n^3 p^3 \ll n^2 p$, so we can remove $o(n^2 p)$ edges to make the graph triangle-free.

Proof. We prove a slightly weaker result, namely that the result is true if $p \gg n^{-1/2} \log n$. The version with $p \gg n^{-1/2}$ can be proved via a stronger formulation of the container lemma (using "fingerprints" as discussed later).

Let $\varepsilon > 0$ be aribtrarily small. Let C be a set of containers for n-vertex triangle-free graphs in Theorem 11.1.1. For every $G \in C$, $e(G) \le \left(\frac{1}{4} + \varepsilon\right) n^2$, so by an application of the Chernoff bound,

$$\mathbb{P}\left(e(G \cap G(n, p)) > \left(\frac{1}{4} + 2\varepsilon\right) n^2 p\right) \le e^{-\Omega_\varepsilon(n^2 p)}$$

Since every triangle-free graph is contained in some $G \in C$, by taking a union bound over C, we see that

$$\begin{aligned}
&\mathbb{P}\left(G(n, p) \text{ has a triangle-free subgraph with } > \left(\frac{1}{4} + 2\varepsilon\right) n^2 p \text{ edges}\right) \\
&\le \sum_{G \in C} \mathbb{P}\left(e(G \cap G(n, p)) > \left(\frac{1}{4} + 2\varepsilon\right) n^2 p\right) \\
&\le |C|\, e^{-\Omega_\varepsilon(n^2 p)} \\
&\le e^{O_\varepsilon(n^{3/2} \log n) - \Omega_\varepsilon(n^2 p)} \\
&= o(1)
\end{aligned}$$

provided that $p \gg n^{-1/2} \log n$. □

11.2 Graph containers

Theorem 11.2.1 (Container theorem for independent sets in graphs)
For every $c > 0$, there exists $\delta > 0$ such that the following holds.

Let $G = (V, E)$ be a graph with average degree d and maximum degree at most cd. There exists a collection $\mathcal{C}$ of subsets of V, with

$$|\mathcal{C}| \leq \binom{|V|}{\leq 2\delta |V| / d}$$

such that

1. Every independent set I of G is contained in some $C \in \mathcal{C}$.
2. $|C| \leq (1 - \delta) |V|$ for every $C \in \mathcal{C}$.

Each $C \in \mathcal{C}$ is called a "container." Every independent set of C is contained in some container.

***Remark* 11.2.2.** The requirement $|C| \leq (1 - \delta) |V|$ looks quite a bit weaker than in Theorem 11.1.1, where each container is only slightly larger than the maximum independent set. In a typical application of the container method, one iteratively applies the (hyper)graph container theorem (e.g., Theorem 11.2.1 and later Theorem 11.3.1) to the subgraphs induced by the slightly smaller containers in the previous iteration. One iterates until the containers are close to their minimum possible size.

For this iterative application of container theorem to work, one usually needs a ***supersaturation*** result, which, roughly speaking, says that every subset of vertices that is slightly larger than the independence number necessarily induces a lot of edges. This property is common to all standard applications of the container method.

The container theorem is proved using

The graph container algorithm (for a fixed given graph G)

Input: a maximal independent set $I \subseteq V$.

Output: a "fingerprint" $S \subseteq I$ of size $\leq 2\delta |V| / d$, and a container $C \supseteq I$ which depends only on S.

Throughout the algorithm, we will maintain a partition $V = A \cup S \cup X$, where

- A, the "available" vertices, initially $A = V$
- S, the current fingerprint, initially $S = \emptyset$
- X, the "excluded" vertices, initially $X = \emptyset$.

The *max-degree order* of $G[A]$ is an ordering of A in by the degree of the vertices in $G[A]$, with the largest first, and breaking ties according to some arbitrary predetermined ordering of V.

While $|X| < \delta |V|$:

1. Let v be the first vertex of $I \cap A$ in the max-degree order on $G[A]$.
2. Add v to S.
3. Add the neighbors of v to X.
4. Add vertices preceding v in the max-degree order on $G[A]$ to X.
5. Remove from A all the new vertices added to $S \cup X$.

Claim: when the algorithm terminates, we obtain a partition $V = A \cup S \cup X$ such that $|X| \geq \delta |V|$ and $|S| \leq 2\delta |V| / d$.

Proof idea: due to the degree hypotheses, in every iteration, at least $\geq d/2$ new vertices are added to X (provided that $d \leq 2\delta |V|$). See Morris' lecture notes for details.

Key facts:

- Two different maximal independent sets $I, I' \subseteq V$ that produce the same fingerprint S in the algorithm necessarily produces the same partition $V = A \cup S \cup X$
- The final set $S \cup A$ contains I (since only vertices not in I are ever moved to I)

Therefore, the total number possibilities for containers $S \cup A$ is at most the number of sets $S \subseteq V$. Since $|S| \leq 2\delta |V| / d$ and $|A \cup S| \leq (1 - \delta) |V|$, this concludes the proof of the graph container lemma.

The fingerprint obtained by the proof actually gives us a stronger consequence that will be important for some applications.

Theorem 11.2.3 (Graph container theorem, with fingerprints)
For every $c > 0$, there exists $\delta > 0$ such that the following holds.

Let $G = (V, E)$ a graph with average degree d and maximum degree at most cd.

Writing $\mathcal{I}$ for the collection of independent sets of G, there exist functions

$$S\colon \mathcal{I} \to 2^V \qquad \text{and} \qquad A\colon 2^V \to 2^V$$

(one only needs to define $A(\cdot)$ on sets in the image of S)
such that, for every $I \in \mathcal{I}$,

- $S(I) \subseteq I \subseteq S(I) \cup A(S(I))$
- $|S(I)| \le 2\delta |V| / d$
- $|S(I) \cup A(S(I))| \le (1 - \delta) |V|$

11.3 Hypergraph container theorem

An independent set in a hypergraph is a subset of vertices containing no edges.

Given an r-uniform hypergraph H and $1 \le \ell < r$, we write

$$\Delta_\ell(H) = \max_{A \subseteq V(H) : |A| = \ell} \text{the number of edges containing } A$$

Theorem 11.3.1 (Container theorem for 3-uniform hypergraph)
For every $c > 0$ there exists $\delta > 0$ such that the following holds.

Let H be a 3-uniform hypergraph with average degree $d \ge \delta^{-1}$ and

$$\Delta_1(H) \le cd \qquad \text{and} \qquad \Delta_2(H) \le c\sqrt{d}.$$

Then there exists a collection $\mathcal{C}$ of subsets of $V(H)$ with

$$|\mathcal{C}| \le \binom{v(H)}{\le v(H)/\sqrt{d}}$$

such that

- Every independent set of H is contained in some $C \in \mathcal{C}$, and
- $|C| \le (1 - \delta) v(H)$ for every $C \in \mathcal{C}$.

Like the graph container theorem, the hypergraph container theorem is proved by designing an algorithm to produce, from an independent set $I \subseteq V(H)$, a fingerprint $S \subseteq I$ and a container $C \supset I$.

The hypergraph container algorithm is more involved compared to the graph container algorithm. In fact, the 3-uniform hypergraph container algorithm calls the graph container algorithm.

Container algorithm for 3-uniform hypergraphs (a very rough sketch):

Throughout the algorithm, we will maintain

- A fingerprint S, initially $S = \emptyset$
- A 3-uniform hypergraph A, initially $A = H$
- A graph G of "forbidden" pairs on $V(H)$, initially $G = \emptyset$

While $|S| \le v(H)/\sqrt{d} - 1$:

- Let u be the first vertex in I in the max-degree order on A
- Add u to S
- Add xy to $E(G)$ whenever $uxy \in E(H)$
- Remove from $V(A)$ the vertex u as well as all vertices proceeding u in the max-degree order on A
- Remove from $V(A)$ every vertex whose degree in G is larger than $c\sqrt{d}$.
- Remove from $E(A)$ every edge that contains an edge of G.

Finally, it is will be the case that either

- We have removed many vertices from $V(A)$
- Or the final graph G has at least $\Omega(\sqrt{d}n)$ edges and has maximum degree $O(\sqrt{d})$, so that we can apply the graph container lemma to G.

In either case, the algorithm produces a container with the desired properties. Again see Morris' lecture notes for details.