Rödl Nibble

Joonleng Tan

1 Hypergraphs

Hypergraphs are a generalisation of graphs. Edges are defined to be a family of distinct subsets of vertices. In hypergraphs, they can connect more than 2 edges together. Rödl's paper[4] deals with r-uniform hypergraphs, which are special hypergraphs where all edges are subsets of r vertices.

2 The Erdös-Hanani conjecture

The Rödl nibble was conceived as a response to the Erdös-Hanani conjecture[2]: for $2 \leq r < k < n$, let $m(n, k, r)$ be the maximal size of a family $\mathcal{F}$ of k-element subsets of $[n]$, with the property that no r vertices lie in more than one $A \in F$. We have

$$m(n, k, r) \leq \frac{\binom{n}{r}}{\binom{k}{r}} \tag{1}$$

The Erdös-Hanani conjecture is simply that

$$\frac{\binom{k}{r}}{\binom{n}{r}} m(n, k, r) \lim_{n \to \infty} 1 \tag{2}$$

If we were to formulate this as a hypergraph problem, let G be the complete r graph. Then $m(n, k, r)$ is the maximal number of disjoint k-cliques that may be "packed" into G. There is a similar problem, the hypergraph cover problem, where $M(n, k, r)$ is the minimal size of a family $\mathcal{F}$ of k-element subsets of $[n]$ with the property that all points lie in at least one $A \in \mathcal{F}$. Both the hypergraph packing problem and the hypergraph cover problem can be solved using an algorithm proposed by Vojtěch Rödl[3], and which came to be known as the Rödl nibble.

3 The Rödl nibble

What follows is a summary of the paper by Rödl[3]. This particular proof is more involved as it tries to draw the connections between the Rödl nibble and the alternative approach to hypergraph proposed by Joel Spencer[5]. A more concise description of the Rödl nibble itself, albeit one that deals with the set cover problem rather than the set packing problem, is given in [1]. It is important to remember that in practice, the method of choosing the random packing is just the random greedy algorithm, and the Rödl nibble is mainly a construct to illustrate ideas with.

We consider the class of r-uniform hypergraphs satisfying $\deg(x) = D(1 + o(1))$ for each vertex $x \in V(H)$ and $\deg(x, y) = o(D)$ for each pair of vertices $x, y \to \infty$. The Rödl nibble involves, at each iteration, selecting at random a set of edges with probability $\frac{\epsilon}{D}$. It is possible to either put an upper bound on the extent to which this set of edges overlap, or the probability that the set of edges selected this way will overlap. After each iteration, the vertices associated with these edges are removed and Rödl shows that, as in the original hypergraph, the degree and co-degrees are also bounded, and the number of vertices in this second hypergraph is concentrated around a predictable mean. This iteration repeats itself until we find

$$|P| \geq \frac{n}{r}(1 - o(1)) \tag{3}$$

3.1 Main theorem: 1.2

The main theorem is that the Rödl nibble will asymptotically give a packing as n grows to infinity:

Let H be an r-uniform hypergraph on a vertex set V of size n. Assume that for each vertex $x \in V(H)$, $\deg(x) = D(1 + o(1))$ and for each pair of vertices, $\deg(x, y) = o(D)$ with $D \to \infty$. $o(1) \to 0$ as $D \to \infty$. Then the random greedy algorithm almost surely gives a packing P such that (3) holds.

3.2 Two lemmas

I quote from the Rödl paper 2 lemmas which are used in the proof. Even though the sketch of the proof provided here will not directly use these 2 lemmas, they are important for understanding how the Rödl nibble works.

3.2.1 Lemma 2.1

Fix $r > 1$, $\epsilon > 0$ and $\gamma > 0$. Then for each $\delta' > 0$ there exists $\delta > 0$ such that the following holds. Suppose that n and D are sufficiently large $(n > n_0(r, \delta', \epsilon, \gamma))$. Let $H = (V(H), E(H))$ be an r-uniform hypergraph with $|V(H)| = n > n_0$ vertices, and let $D > D_0$. Suppose that:

(i) for each vertex $x : \deg_H(x) < \gamma D$
(ii) for all but at most δn vertices $x : \deg_H(x) = D(1 \pm \delta)$
(iii) for each pair of vertices $x, y : \deg_H(x, y) < \delta D$

Let **N** be a random subfamily of $E(H)$ such that each edge is chosen into **N** independently with probability $\frac{\epsilon}{D}$ then if $\mathbf{W} = V(H) \setminus \cup \mathbf{N}$ and **H***$= H \setminus \{\cup \mathbf{N}\}$, the following holds almost surely as $n \to \infty$ and $D \to \infty$):

(iv) $|\mathbf{N}| = \frac{\epsilon n}{r}(1 \pm \delta')$,
(v) $|\mathbf{W}| = ne^{-\epsilon}(1 \pm \delta')$,
(vi) $\deg_{\mathbf{H}*}(x) = De^{-\epsilon(r-1)}(1 \pm \delta')$ for all but at most δ' $|\mathbf{W}|$ vertices of $\mathbf{H}*$.

This lemma guarantees that after selecting edges out of a hypergraph at random with probability $\frac{\epsilon}{D}$, the hypergraph resulting from the remaining vertices has properties very similar to (i) and (ii).

3.2.2 Lemma 2.2

Fix $\epsilon > 0$ and δ, $0 < \delta < \frac{1}{4}$. Let $H = (V(H), E(H))$ be an r-uniform hypergraph on a set of n vertices such that for each vertex $x, \deg(x) = D(1 \pm \delta)$. Let **N** be a random subset of edges defined by

$$P(e \in \mathbf{N}) = \frac{\epsilon}{D}, \forall e \in E(H) \tag{4}$$

Then the number of collisions (ie pairs of edges a, b such that $a \cap b \neq \emptyset$) in **N** is almost surely (as $n \to \infty$ and $D \to \infty$) bounded from above by $n\epsilon^2$.

This shows that for the remaining hypergraph, the codegree condition will still hold in some form. This also shows that for the edges that are removed during the iteration, there is a bounded degree of overlap that the proof of the main theorem will attempt to deal with.

3.3 Proof of main theorem (sketch)

I will provide a sketch of the proof: the details are in the main paper. Given α, we define:

$$\epsilon = \alpha(\frac{\alpha}{3})^{2r+1} \quad (5)$$

$$t_1 = \frac{-1}{\epsilon}\ln[\alpha(1-5\alpha)] \quad (6)$$

This proof will show that we can get a packing of size greater than $\frac{n}{r}(1-\alpha)$ almost surely as $D \to \infty$. There are 2 processes to consider. Process $\mathcal{R}$ shows us how the "nibbles", or $\mathbf{N}_i$ are generated. Notice that 2 edges in different $\mathbf{N}_i$s will not overlap, but within each $\mathbf{N}_i$ there is a possibility that the edges will overlap. The union of nibbles, $\mathbf{P}^{\mathcal{R}}$ does not constitute a packing, so a modification has to be done and is described in Process $\mathcal{S}$.

Algorithm 1 Generative process $\mathcal{R}$

for i = 1 to t_1 **do**
  $D_i \leftarrow D^{-e(i-1)(r-1)}$
  $p_i \leftarrow \frac{\epsilon}{D_i}$
  for each $e \in \mathbf{H} - \cup_{j=1}^{i}\mathbf{B}_j$ **do**
    $\mathbf{B}_i = \mathbf{B}_i \cup e$ with probability p_i
  end for
  if $i = 1$ **then**
    $V(\mathbf{H}_1) = V(\mathbf{H})$
    $E(\mathbf{H}_1) = E(\mathbf{H})$
    $\mathbf{N}_1 = \mathbf{B}_1$
    $\mathbf{I}_1 = \emptyset$
  else
    $V(\mathbf{H}_i) = V(\mathbf{H}_{i-1}) - \cup_{e \in N_{i-1}} e$
    $E(\mathbf{H}_i) = \{e \subset E(\mathbf{H}_{i-1}); e \in V(\mathbf{H}_i)\ \}$
    $\mathbf{N}_i = \mathbf{B}_i \cap E(\mathbf{H}_i)$
    $\mathbf{I}_i = \mathbf{B}_i - \mathbf{N}_i$
  end if
  let $\prec_i$ be a permutation on $\mathbf{B}_i$ selected with uniform probability
end for
$\mathbf{B}_{t_1+1} = E(\mathbf{B}) \setminus (\cup_{j=1}^{t_1}\mathbf{B}_j)$
$\mathbf{N}_{t_1+1} = \emptyset$
$\mathbf{I}_{t_1+1} = \mathbf{B}_{t_1+1}$
$\mathbf{P}^{\mathcal{R}} \leftarrow \cup_{j=1}^{t_1+1}\mathbf{N}_j$

Algorithm 2 Generative process $\mathcal{S}$

$\mathbf{P}^{\mathcal{S}} = \emptyset$
for each $e \in E(\mathbf{H})$, from least edge to largest edge in the ordering given by $\{\prec_1, \ldots, \prec_{t_1+1}\}$ **do**
  if $e \cap (\cup_{a \in P^{\mathcal{S}}} a) = \emptyset$ **then**
    $\mathbf{P}^{\mathcal{S}} \leftarrow \mathbf{P}^{\mathcal{S}} \cup e$
  end if
end for
for i = 1 to t_1 **do**
  $\mathbf{S}_i = \mathbf{P}^{\mathcal{S}} \cup \mathbf{B}_i$
  $\mathbf{R}_i = \mathbf{B}_i - \mathbf{S}_i$
end for

From process $\mathcal{S}$, we have a packing. This approach of randomly ordering the edges and choosing them in ascending order is similar to Joel Spencer's approach and the interesting feature of this particular proof is how Rödl shows that the greedy approach and the nibble approach are similar.

After this, we define some events. The paper provides more technical details on the probability space. Let δ_i be suitable bounds pertaining to each tranche of edges.

$\mathcal{A}_i = \{\prec_1, \ldots, \prec_{t_1+1}$, such that
$|\mathbf{N}_i| = \frac{\epsilon n}{r}.e^{\epsilon(i-1)}(1 \pm \delta_i)^i$,

$V(\mathbf{H}_{i+1}) = ne^{\epsilon i}(1 \pm \delta_i)^i\}$
and all but at most $\delta_i \mid V(\mathbf{H}_{i+1}) \mid$ vertices satisfy $\deg H_{i+1}(x) = De^{\epsilon i(r-1)}(1 \pm \delta_i)\}$

$\mathcal{C}_i = \{\prec_1, \ldots, \prec_{t_1+1}$, such that
$|\mathbf{N}_i \Delta \mathbf{S}_i \, |\leq 8n\epsilon^2 e^{2r\epsilon(i-1)}\}$

Here Δ refers to the symmetric difference of sets. From Lemma 2.1, we get that

$$P(\mathcal{A}_1) = 1 - o(1) \tag{7}$$
$$P(\mathcal{A}_i \mid \mathcal{A}_{i-1}) = 1 - o(1) \tag{8}$$

These imply that $P(\mathcal{A}_i) = P(\mathcal{A}_i \mid \mathcal{A}_{i-1})P(\mathcal{A}_{i-1}) = 1 - o(1)$. Furthermore, setting $\mathcal{A} = \mathcal{A}_1 \cap \mathcal{A}_2 \cap \cdots \cap \mathcal{A}_{t_1}$ give us

$$P(\mathcal{A}) = 1 - o(1) \tag{9}$$

Where $o(1) \to 0$ as $D \to \infty$.

3.4 Two more lemmas

3.4.1 Claim 2.3

$$P(\mathcal{C}_1) = 1 - o(1) \tag{10}$$
$$P(\mathcal{C}_i \mid \mathcal{A}_1 \cap \cdots \cap \mathcal{A}_{i-1} \cap \mathcal{C}_1 \cap \cdots \cap \mathcal{C}_{i-1}) = 1 - o(1) \tag{11}$$

For $i = [t_1]$. Where $o(1) \to 0$ as $D \to \infty$.

3.4.2 Claim 2.4

Let $\mathcal{C} = \mathcal{C}_1 \cap \mathcal{C}_2 \cap \cdots \cap \mathcal{C}_{t_1}$.

$\mathcal{A} \cap \mathcal{C} \subseteq \{\prec_1, \ldots, \prec_{t_1+1}$ such that $\mid P^{\mathcal{S}} \mid \geq \frac{n}{r}(1-\alpha)\}$

Translation: $\mathcal{A} \cap \mathcal{C}$ means that the packing found by $\mathcal{S}$ satisfies the bounds that we are looking for. The details of these claims are in the paper but are not for the faint of heart.

What follows are some technical details on how to translate the events from the probability space over $\prec_i$ to the probability space over σ of all possible permutations of V(**H**). Then we prove that

$$P(\mathcal{A} \cap \mathcal{C}) = 1 - o(1) \tag{12}$$

by induction. Base case is (7) and (10). Then assuming that $P(\mathcal{C}_j) = 1 - o(1)$ for $j \in [i-1]$, we show that

$$\begin{aligned} P(\mathcal{C}_i) \geq & P(\mathcal{C}_i \cap \mathcal{A}_1 \cap \cdots \cap \mathcal{A}_{i-1} \cap \mathcal{C}_1 \cap \cdots \cap \mathcal{C}_{i-1}) \\ = & P(\mathcal{C}_i \mid \mathcal{A}_1 \cap \cdots \cap \mathcal{A}_{i-1} \cap \mathcal{C}_1 \cap \cdots \cap \mathcal{C}_{i-1}) \\ & P(\mathcal{A}_1 \cap \cdots \cap \mathcal{A}_{i-1} \cap \mathcal{C}_1 \cap \cdots \cap \mathcal{C}_{i-1}) \\ = & 1 - o(1) \end{aligned}$$

The last equality follows from (9) and (11). This proves (12). QED for main theorem.

3.5 Proposition 1.3

The other main result of the paper is the proposition which concerns itself with how many edges are processed by the algorithm before the packing of the desired size is reached.

Let $r \geq 2$ be an integer and $\alpha > 0$ be a sufficiently small positive real number. Then there exists $\delta = \delta(\alpha, r) > 0$ and $n_0, D_0 > 0$ such that for each $n > n_0, D > D_0$ and for any r-uniform hypergraph $H = (V(H), E(H)), \mid V(H) \mid = n$, satisfying the assumptions of Theorem 1.2, the running time T_α (ie the number of processed edges by the algorithm) of the random greedy algorithm which stops after yielding a packing of size $\lceil \frac{n}{r}(1-\alpha) \rceil$ almost surely satisfies the following inequalities:

$$\frac{1-5\alpha}{[\alpha(1+5\alpha)]^{r-1}} \frac{n}{r(r-1)} \leq T_\alpha \leq \frac{1+5\alpha}{[\alpha(1-5\alpha)]^{r-1}} \frac{n}{r(r-1)} \tag{13}$$

In particular, $T_\alpha \sim (\frac{1}{\alpha})^{r-1} \frac{n}{r(r-1)}$ as $\alpha \to 0^+$

References

[1] Noga Alon and Joel Spencer. Asymptotic packing via a branching process. *The Probabilistic Method, 2nd ed.*, John Wiley, 2000.

[2] Paul Erdös and Haim Hanani. On a limit theorem in combinatorial analysis. *Publ. Math. Debrecen*, 10:10–13, 1963.

[3] Rödl, Vojtěch and Thoma, Luboš. Asymptotic packing and the random greedy algorithm. *Random Structures & Algorithms*, 8(3):161–177, 1996.

[4] Rödl, Vojtěch. On a Packing and Covering Problem. *European Journal of Combinatorics*, 5:69–78, 1985.

[5] Joel Spencer. Asymptotic packing via a branching process. *Random Structures & Algorithms*, 7(2):167–172, 1995.