Serious Combo 2

Po-Shen Loh

June 2013

1 Definitions

- An r-set is a set of size r.
- The set $\{1, 2, \ldots, n\}$ is often denoted $[n]$.
- Given a set S, its power set 2^S is the collection of all subsets of S.
- A hypergraph is a set of vertices V, together with a collection $E \subset 2^V$ of subsets of V. The members of E are called hyperedges, or just edges for short.
- An r-uniform hypergraph is a hypergraph in which all edges have are r-subsets of V. Note that a 2-uniform hypergraph is just a graph.
- If $U \subset V$, the sub-hypergraph $H[U]$ of H induced by U is the hypergraph with vertex set U, together with all edges of the original H which were completely contained within U.
- An independent set in a hypergraph is a collection of vertices which induces no edges.

2 Problems and famous results

1. (From Peter Winkler.) The 53 MOPpers were divided into 7 teams for Team Contest 1. They were then divided into 6 teams for Team Contest 2. Prove that there must be a MOPper for whom the size of her team in Contest 2 was strictly larger than the size of her team in Contest 1.

   Solution: In Contest 1, suppose the team breakdown was $s_1 + \cdots + s_7 = 60$. Then in the i-th team, with s_i people, say that each person did $\frac{1}{s_i}$ of the work. Similarly, in Contest 2, account equally for the work within each team, giving scores of $\frac{1}{s'_i}$.

   However, the total amount of work done by all people in Contest 1 was then exactly 7, and the total amount of work done by all people in Contest 2 was exactly 6. So somebody must have done strictly less work in Contest 2. That person saw
   $$\frac{1}{s'_i} < \frac{1}{s_i},$$
   i.e., the size of that person's team on Contest 2 was strictly larger than her team size on Contest 1.

2. (Sauer-Shelah.) A family $\mathcal{F}$ shatters a set A if for every $B \subset A$, there is $F \in \mathcal{F}$ such that $F \cap A = B$. Prove that if $\mathcal{F} \subset 2^{[n]}$ and $|\mathcal{F}| > \binom{n}{0} + \cdots + \binom{n}{k}$, then there is a set $A \subset [n]$ of size $k+1$ such that $\mathcal{F}$ shatters A.

   Solution: Compressions: show that if a set family has no shattered subset of size t, then after compressing, it still has no shattered subset of size t. To visualize this, a shattered subset S of size t is specified by a collection of t elementary basis vectors in the hypercube, and the plane P determined by the 2^t vectors in their span. Now, consider the orthogonal complement of P, and imagine partitioning the entire hypercube into $|P|$ fibers, each of which is running orthogonally into P. The shattering means that each of the fibers contains some point in the set family. Now a compression is a wind which blows in one dimension. If the wind blows orthogonally to P, and S is shattered, then it is clear that S was also shattered before. Also, if the wind blows parallel to P, and S is now shattered, then it is also clear that S was shattered before. Therefore, we may assume that all compressions have been completed, and now if a set is in the family, all of its subsets are as well. This forces there to be a set of size $k+1$.

3. (Erdős-Ko-Rado.) Let $\mathcal{F}$ be a collection of distinct r-subsets of $[n]$, with the property that every pair of subsets $A, B \in \mathcal{F}$ intersects. Prove that $|\mathcal{F}| \le \binom{n-1}{r-1}$ whenever $n \ge 2r$.

4. (Putnam 1956/A7.) Call an integer $0 \le r \le 1000$ *strange* if the number of r-subsets of $\{1, \ldots, 1000\}$ is odd. Prove that the number of strange integers is a power of 2. (To make this an AIME problem: how many strange integers are there?)

   Solution: Lucas's theorem: it is going to be 2 raised to the power which is equal to the number of 1's in the binary expansion of 1000.

5. (Turán.) Let G be an n-vertex graph with average degree d. Then it contains an independent set of size at least $\frac{n}{d+1}$, and this is tight.

6. Let $\mathcal{F}$ be a collection of r-subsets of $[n]$, and let $t = |\mathcal{F}|/n$. Then there is always a set $S \subset [n]$ of size at least $n/(4t^{\frac{1}{r-1}})$, which does not completely contain any member of $\mathcal{F}$.

7. (Diestel 2.20.) Let G be a graph, and let α be the size of its largest independent set. Prove that the vertices of G can be covered by α disjoint subgraphs, each either a cycle, a K_2, or a K_1.

   Solution: Take a longest path, and say that it ends at v. Make a cycle by taking the edge from v back to its earliest neighbor along the path. Delete this cycle from the graph. Importantly, v is not adjacent to any vertex in the rest of the graph! So, in the remainder, the independence number is at most $\alpha - 1$, and we may apply induction. Note that K_2 arises when the longest path has only one edge, in which case we can't close a cycle, and K_1 arises when the longest path is a single vertex.

8. (Folklore.) The chromatic number χ of a hypergraph H is the minimum integer k such that it is possible to assign a color from $[k]$ to each vertex of H, with no edge having all of its vertices in the same color. (Assume that there are no edges of size 1.) Prove that if $\chi > 2$, then H must have two edges which intersect in exactly one vertex.

9. (L.) An r-tree is an r-uniform hypergraph created in the following way: starting with a single hyperedge of size r, repeatedly add new hyperedges by selecting one existing vertex v, and adding $r-1$ new vertices, together with a new hyperedge through v and the $r-1$ new vertices. Let T be an arbitrary r-tree with t edges. Observe that T will always have exactly $1 + (r-1)t$ vertices. Prove that every r-uniform hypergraph H with chromatic number $\chi > t$ must contain T as a sub-hypergraph.

10. Let G be a graph in which all vertices have nonzero degree. Prove that its vertices can be partitioned into two sets $V_1 \cup V_2$ such that the number of edges going between the V_i is at least $\frac{m}{2} + \frac{n}{6}$. Is this tight?

    Solution: Analyze the greedy partitioning algorithm, which puts each new vertex onto the side which maximizes the number of crossing edges to already-placed vertices. Observe that we gain $\frac{1}{2}$ per vertex which had odd back-degree. But if we take a random ordering of the vertices, then the expected number of vertices with odd back-degree is at least $m/3$, with sharpness on vertices of degree 2 (since their back-degree is 0, 1, or 2).

11. (MOP 2007/7/1.) In a 100×100 array, each of the numbers $1, 2, \ldots, 100$ appears exactly 100 times. Show that there is a row or a column in the array with at least 10 distinct numbers.

Solution: Let $n = 100$. Choose a random row or column ($2n$ choices). Let X be the number of distinct entries in it. Now $X = \sum I_i$, where each I_i is the indicator variable of i appearing (possibly more than once) in our random row or column. Clearly, each $\mathbb{E}[I_i] = \mathbb{P}[I_i \geq 1]$. To lower-bound this, observe that the worst-case is if all n appearances of i are in some $\sqrt{n} \times \sqrt{n}$ submatrix, which gives $\mathbb{P}[I_i \geq 1] \geq 2\sqrt{n}/(2n) = 1/\sqrt{n}$. Hence by linearity, $\mathbb{E}[X] \geq \sqrt{n}$.

12. (Russia, 1999.) In a class, each boy is friends with at least one girl. Show that there exists a group of at least half of the students, such that each boy in the group is friends with an odd number of the girls in the group.

    Solution: Choose girls independently with probability 1/2, and then let the set of boys be all of those who have an odd number of friends in the girl group. Let X be the number of boys and girls selected, and break this into the sum of indicators. For each girl, obviously the indicator adds 1/2 to the sum. For each boy, the probability that he joins is precisely the probability that $\text{Bin}[k, 1/2]$ is odd, where k was the number of girls he knew. To see that this probability is 1/2, note that it is the parity of the sum of k independent coin flips. In particular, the final flip independently flips or retains the final parity, hence odd with probability 1/2.

13. (IMO Shortlist 1999/C4.) Let A be any set of n residues mod n^2. Show that there is a set B of n residues mod n^2 such that at least half of the residues mod n^2 can be written as $a + b$ with $a \in A$ and $b \in B$.

    Solution: Make n independent uniformly random choices from the n^2 residues, and collect them into a set B. Note that since we use independence, this final set may have size $< n$. But if we still have $A + B$ occupying at least half of the residues, then this is okay (we could arbitrarily augment B to have the full size n).

    Let X be the number of residues achievable as $a + b$. For each potential residue i, there are exactly n ways to choose some b for which $A + b \ni i$, since $|A| = n$. Therefore, the probability that a given residue i appears in $A + B$ is precisely $1 - \left(1 - \frac{n}{n^2}\right)^n$. Then $\mathbb{E}[X]$ is exactly n^2 times that, because there are n^2 total residues. Hence it suffices to show that $1 - \left(1 - \frac{n}{n^2}\right)^n \geq 1/2$. But this follows from the bound $1 - \frac{1}{n} \leq e^{-1/n}$, using $e \approx 2.718$.

14. (Gallai, Hasse, Roy, Vitaver.) Let D be a directed graph, and let χ be the chromatic number of its underlying undirected graph. Show that D has a directed path of at least χ vertices.

    Solution: Take a maximal acyclic subgraph, and use it to define level sets, by coloring each vertex by the length of the longest directed path that ends at that vertex.

15. (MOP 2010, harder variation.) Let G be a graph with average degree d. Prove that for every $k \leq d$, there is a K_{k+1}-free induced subgraph on at least $\frac{kn}{d+1}$ vertices.

    Solution: Randomly permute the vertices. Use the greedy algorithm, taking each vertex if it can be added without making any K_{k+1}. Observe that we will actually take every vertex v with the property that the permutation induced on $\{v\} \cup N(v)$ has v in position $1 \ldots k$. (We might also take more.) This is because v would only have degree at most $k - 1$ back to the previously selected guys, making at most a K_k. Now the expected size of the selected set is at least

    $$\sum_v \frac{k}{d_v + 1} \geq n \cdot \frac{k}{d+1}.$$

16. (Oddtown.) Let $\mathcal{F}$ be a collection of distinct subsets of $2^{[n]}$ such that every $A \in \mathcal{F}$ has size which is nonzero modulo 2, but every pair of distinct $A, B \in \mathcal{F}$ has intersection size which is zero modulo 2. Prove that $|\mathcal{F}| \leq n$.

17. (Open.) What if 2 is replaced by 6?

18. (Russia 2006, final problem.) A group of pioneers has arrived to summer camp. Each pioneer has at least 50 and at most 100 friends among the others. Prove that one can distribute field caps of 1331 colors among the pioneers so that the friends of each pioneer have caps of at least 20 colors.

Solution: We prove a much stronger bound. Let $C = 49$ be the total number of colors, and give each person an independent, uniformly random color. We will apply the Lovász Local Lemma. For each vertex, let B_v be the event that $N(v)$ receives 19 or fewer colors. We have:

$$\mathbb{P}[B_v] \leq \binom{C}{19}\left(\frac{19}{C}\right)^{50}.$$

Now we build the dependency graph. Consider all vertices in $N(v)$, together with those adjacent to $N(v)$. Connect each vertex to v in the dependency graph. It is clear that all other vertices do not need to be connected to v. Hence the dependency is below 10^4. We therefore have a solution when

$$e \cdot \binom{C}{19}\left(\frac{19}{C}\right)^{50} \cdot 10^4 < 1,$$

Note in particular that $\binom{C}{19} \leq \left(\frac{eC}{19}\right)^{19}$, so it suffices to have:

$$\begin{aligned} e\left(\frac{eC}{19}\right)^{19}\left(\frac{19}{C}\right)^{50} 10^4 &< 1 \\ e^{20} 19^{31} 10^4 &< C^{31} \\ 48.75 \approx 19 e^{\frac{20}{31}} 10^{\frac{4}{31}} &< C. \end{aligned}$$