Probabilistic Method

The classic approaches to existence results are either inductive, constructive or *reductio ad absurdum*. Probabilistic Method (PM) is another alternative. In existence proofs, we are given a set of objects and we want to show the existence of a subset that has a desired property. The usual approach in PM is to pick an object from the set randomly (according to some suitable distribution) and show that the probability of picking an object with the desired property is non-zero, thus implying its existence; occasionally, it is better to show that the probability of picking an object not having the desired property is strictly smaller than one. If our universe set is countable, then one can say that there is no necessity to introduce probability, as we can enumerate all the elements and check them if they have the desired property or not; though this approach is constructive, the size of the universe sets is usually huge and so the approach is not computationally feasible; on the contrary, PM-based proofs are usually slick. The use of basic tools (expectation, variance, Markov's inequality, Chebyshev's inequality) and not so basic tools (Lovász Local Lemma, Martingales, Random Walks) from probability have led to various breakthroughs. In this lecture, we present some of these results.[1]

1 Tournaments

A **tournament** T_n is an orientation of K_n, i.e., to each edge (i, j) we assign a direction $i \to j$ or $j \to i$, where "$i \to j$" is to be interpreted as i defeats j in a match; there cannot be any ties. Formally, $T = (V, E)$, where $V = [n]$ and E is the edges of K_n with their orientations.

Given a k, a tournament T is said to have **property** S_k if for every subset of k players, there is a remaining player who defeats them all. E.g., the triangle with edges $\{(1, 2), (2, 3), (3, 1)\}$ has S_1 but not S_2; in fact it cannot have S_2. Can a graph on four vertices have S_2? No! How about five vertices? No, again! How many vertices should we have to get a tournament with property S_2, or in general S_k? This problem was raised by Schütte and resolved by Erdős (1963). PM gives us a sufficient condition.

Theorem 1 (Erdös (1963)). *Given a k, if*

$$\binom{n}{k}(1 - 2^{-k})^{n-k} < 1$$

then there exists a tournament T_n with property S_k.

Proof.

1. The idea is that for n sufficiently large a random tournament on K_n is likely to have property S_k. What is a random tournament? For every edge (i, j) in K_n with probability half we either orient it $i \to j$ or $j \to i$. Thus all the $2^{\binom{n}{2}}$ tournaments are equally likely.

2. What is the probability that T_n does not have the property S_k? For $S \in \binom{[n]}{k}$, let A_S be the event that S does not have a winner in $V \setminus S$. Then the desired probability is that there exists an S that has no winner, i.e., at least one of the events A_S occurs, which is equal to

$$\Pr\Big(\bigcup_{S \in \binom{[n]}{k}} A_S\Big) \le \sum_{S \in \binom{[n]}{k}} \Pr(A_S). \tag{1}$$

[1] The earliest use of PM was by Szele (1943) in showing the existence of Hamiltonian paths in tournaments, however, it was Erdős who revealed its true potential by applying it extensively.

3. What is $\Pr(A_S)$? The probability that $v \in V \setminus S$ is a winner for S is 2^{-k}, i.e., it has a directed edge to all the vertices in S. Thus the probability that v is not a winner is $1-2^{-k}$. Thus $\Pr(A_S) = (1-2^{-k})^{n-k}$, i.e., no vertex in $v \in V \setminus S$ is a winner w.r.t. S and since all the edges were oriented independently. Plugging this upper bound in (1), we obtain that if

$$\binom{n}{k}(1-2^{-k})^{n-k} < 1$$

then with positive probability no event A_S occurs, and hence there is a tournament T_n that has property S_k.

Q.E.D.

In the proof above, we used the simple observation that for n events $A_1, \ldots, A_n$, the probability that at least one of them occurs satisfies $\Pr(\cup_{i=1}^n A_i) \le \sum_{i=1}^n \Pr(A_i)$.

Suppose we want to know if there is a tournament T_n with a linear order on the players, i.e., an ordering of the players $v_1, \ldots, v_n$ such that $v_1 \to v_2 \to \cdots \to v_{n-1} \to v_n$? Clearly, such a linear order is equivalent to having a Hamiltonian path in T_n. PM can not only tell us whether there is a tournament that has a Hamiltonian path, but can also be used to exhibit tournaments with a large number of Hamiltonian paths, which is a more quantitative result. Szele's result was a result in this direction, and combines PP with PM. The fundamental results that it uses are the linearity of expectation: If $X_1, \ldots, X_n$ are random variables, then

$$\mathtt{E}[X := c_1X_1 + \cdots + c_nX_n] = \sum_{i=1}^{n} c_i\mathtt{E}[X_i]; \tag{2}$$

and the fact that there are points σ_1, σ' in the probability space such that $X(\sigma) \ge \mathtt{E}[X]$ and $X\sigma' \le \mathtt{E}[X]$, which is equivalent to the PP.

Theorem 2 (Szele (1943)). *There is a tournament T_n with $n!/2^{n-1}$ Hamiltonian paths.*

Proof. A Hamiltonian path in T_n is a permutation σ of the vertices such that $(\sigma(i), \sigma(i+1))$, $i = 1, \ldots, n$, is an edge in T_n. Given an arbitrary permutation σ, let $X_\sigma : S_n \to \{0, 1\}$ be the indicator variable telling us whether σ is a Hamiltonian path or not; thus $\mathtt{E}[X_\sigma] = \mathtt{Pr}\{\sigma \text{ is a Hamiltonian path}\} = 1/2^{n-1}$. Then for $X := \sum_{\sigma \in S_n} X_\sigma$, $\mathtt{E}[X]$ is the expected number of Hamiltonian paths in T_n. From (2) it follows that

$$\mathtt{E}[X] = \sum_{\sigma \in S_n} \mathtt{E}[X_\sigma] = \frac{n!}{2^{n-1}}.$$

From PP we know that there is a σ such that $X_\sigma \ge \mathtt{E}[X]$. **Q.E.D.**

2 Chromatic Number and Girth

Chromatic Number of a graph G, $\chi(G)$, is the minimum number of colors required to color G s.t. no two adjacent vertices get the same color. A large value of $\chi(G)$ would seem to suggest that the graph is highly connected, and in particular, that the graph contains a complete graph, which would imply it has triangles. However, this is not the case.[2] Blanche Descartes had already shown that there are graphs with arbitrarily high chromatic number but minimum length of cycles at least four. This latter quantity is called the **girth** of G, $\gamma(G)$. Let's see the construction.

We will construct a sequence of graphs $G_3, G_4, \ldots$ s.t. $\chi(G_n) = n$, but G_n has not triangles. Let $G_3 := C_5$, then clearly $\chi(G_3) = 3$. Let $G_n = (V_n, E_n)$. We construct G_{n+1} from G_n as follows: Let V' be a copy of the vertices of V_n s.t. they are in bijection, $v \in V \to v' \in V'$; connect each vertex $v' \in V'$ to all the neighbors of v in G_n; introduce a new vertex z and connect it to all the vertices in V'; then $V_{n+1} := V_n \cup V' \cup \{z\}$ and $E_{n+1} := E_n$ plus all the edges introduced in the last two steps of the construction.

Claim:

[2] On the contrary, Hadwiger's conjecture states that this is almost the case, i.e., if the chromatic number of the graph is n then the graph contains K_n as a minor.

1. $\chi(G_{n+1}) = n+1$. That $n+1$ colors suffice is easy to see: every vertex $v' \in V'$ can take the same color as its corresponding vertex $v \in V_n$, and z takes a color different from all the n colors used in V'. Why are $n+1$ colors necessary? Considering a coloring of G_n using n colors; let $C_1, \ldots, C_n$ be the partition of V induced by the colors. We claim that within each C_i there must be a vertex v_i that is neighbor to vertices from all the remaining color classes; if not, then for every $v \in C_i$ we can color it with one of the remaining colors and hence we can color G_n with fewer than n colors, which is a contradiction. Therefore, the vertex v_i' corresponding to v_i has to have the same color as v_i. Hence V' needs n colors, and consequently z needs a new colors, which implies that G_{n+1} needs $n+1$ colors.

2. $\gamma(G) = 4$. The cycles already present in G_n have length at least 4. The new cycles are obtained by the edges connecting V_n to V' and the edges connecting V' to z. These new cycles have length at least 4: since to connect $v, w \in V_n$ we need the three edges $v - v' - z - w' - w$; the best we can do is when $w = v$; or in other words, we need one edge to enter V' and two edges to exit from it.

But can we do better than four? Can there be graphs whose coloring and girth both are large?

Theorem 3 (Erdős (1959)). *For every $k > 2$, there exists a graph G such that $\chi(G) > k$ and $\gamma(G) > k$.*

Idea: To construct a probability space and show that the probability that $\chi(G) \le k$ is $< 1/2$ or the probability that $\gamma(G) \le k$ is $< 1/2$. Thus there must be a graph for which both properties are at least k. This time, however, the distribution will not be uniform.

Step 1: Our first step is actually to replace $\chi(G)$ with another parameter $\iota(G)$, the independence number of G. Intuitively, χ and ι are connected, but they are connected in an inverse manner. Any χ coloring of V partitions V into $P_1, \ldots, P_\chi$ sets, where vertices of the same color are in the same partition. Moreover, vertices in the same partition are independent of each other, thus $|P_j| < \iota$. Since P_j's form a partition of V it follows that $|P_1| + |P_2| + \cdots + |P_\chi| = n$. Substituting the upper bound on $|P_j|$'s we obtain

$$\chi\iota \ge n. \tag{3}$$

The probability that χ is small thus implies that ι is large. In fact, we will show that $\texttt{Pr}\{\iota > cn\}$ is small, for some fraction c.

Step 2: What is the probability space? It is the set $\mathcal{G}(n,p)$ of graphs on n vertices where the individual edges appear with probability p, independent of each other; p will be appropriately chosen later. So, the probability that K_n is picked is $p^{\binom{n}{2}}$; in general, a graph with m edges is picked with probability $p^m(1-p)^{\binom{n}{2}-m}$.

Step 3: What is the probability that a $G \in \mathcal{G}(n,p)$ has $\iota(G) \ge r$? What is the probability that $R \subseteq V$, $|R| = r$, is an independent set? It is $(1-p)^{\binom{r}{2}}$. If A_R is the event that R is an independent set, then

$$\Pr(\iota(G) \ge r) = \Pr(\bigcup_{R \in \binom{V}{r}} A_R) \le \sum_{R \in \binom{V}{r}} \Pr(A_R) = \binom{n}{r}(1-p)^{\binom{r}{2}} \le (ne^{-p(r-1)/2})^r,$$

where the last step follows from the fact that for all p, $(1-p) \le e^{-p}$. Recall that we wanted to show that the probability that ι is greater than a certain fraction of n is smaller than half. In particular, choosing $r = n/2k$ in the inequality above we get that there exists N_1, s.t. for all $n > N_1$

$$\Pr(\iota \ge \frac{n}{2k}) < \frac{1}{2}. \tag{4}$$

Step 4: We now derive a similar result for $\gamma(G)$. We want to show that there exists an N_2 s.t. for all $n \ge N_2$, $\Pr(\gamma(G) \le k)$ is small. Let X be the random variable that counts the number of cycles with length $\le k$. We will in fact show that $\Pr(X \ge s)$ is small, i.e., we cannot have too many cycles with length $\le k$. To get this result, we use yet another fundamental result from probability, namely Markov's inequality.

Step 5: What is the probability that a given subset $C \subseteq V$ forms a cycle of length j? Every permutation of $c_1, \ldots, c_j$ can be thought of as a cycle. However, some permutations give rise to the same cycle: all cyclic shifts of a permutation gives the same cycle, and since the graph is undirected, the permutation and its reverse are the same cycle. Thus the number of cycles of length j is $j!/(2j) = (j-1)!/2$. Thus the total number of cycles of length j in G are $\binom{n}{j}(j-1)!/2$; every such cycle appears with probability p^j.

Step 6: Recall that X is the random variable that counts the number of cycles with length $\leq k$. With each cycle C, we associate the indicator random variable X_C that is one iff C appears in G; thus $\mathtt{E}[X_C] = p^{|C|}$. Clearly, $X = \sum_{C||C|\leq k} X_C$. Since X is a positive random variable, from Markov's inequality we obtain

$$\mathtt{Pr}\{X \geq s\} \leq \frac{\mathtt{E}[X]}{s}.$$

Moreover, from the linearity of expectation we have

$$\mathtt{Pr}\{X \geq s\} \leq \frac{1}{s}\sum_{C||C|\leq k} \mathtt{E}[X_C] = \frac{1}{s}\sum_{j=3}^{k}\binom{n}{j}(j-1)!p^j/2 \leq \frac{1}{s}\frac{1}{2}\sum_{j}(np)^j \leq (k-2)(np)^k/2.$$

Thus

$$\mathtt{Pr}\{X \geq n/2\} \leq k\frac{(np)^k}{n}.$$

An appropriate choice of p, yields that the RHS is smaller than $1/2$. Thus there exists N_2 s.t. for all $n \geq N_2$

$$\mathtt{Pr}\{X(G) \geq n/2\} < \frac{1}{2}. \tag{5}$$

Step 7: Let $N \geq \max N_1, N_2$. Then from (4) and (5) it follows that there is a graph H on N vertices such that

$$\iota(H) < \frac{N}{2k} \text{ and } X(H) < \frac{N}{2}.$$

How small is $\chi(H)$? From (3) it follows that $\chi(H) \geq n/\iota(H) > 2k$. So H has a large chromatic number, but not girth as it has cycles with length $\leq k$. We construct a new graph G from H by deleting one vertex from these cycles thus breaking all the cycles of length $\leq k$. The resulting graph G is not empty as we have deleted at most $N/2$ vertices, and clearly $\gamma(G) > k$. What about $\iota(G)$? Well, $\iota(G) \leq \iota(H)$, as we have removed vertices from, say, the simple cycles, and that does not increase ι. Again applying (3) to G we get

$$\chi(G) \geq \frac{N}{2\iota(G)} \geq \frac{N}{2\iota(H)} > k.$$

Thus G has both the desired properties.

3 Lovász Local Lemma

Given an n-uniform hypergraph H, we say H is k-colorable if there is a k-coloring of its vertices such that no edge is monochromatic. When can we say that H is, let's say, 2-colorable? Let's apply the PM to this problem when H is n-uniform. If we color the set of vertices randomly with 2-colors, where we can pick any color with equal probability, then the probability that an edge is monochromatic is 2^{1-n}, since all the n vertices get the same color with prob 2^{-n} and there are two colors to choose from. Let B_e be the event that e is monochromatic. Thus probability that at least one of B_e occurs is

$$\Pr(\bigcup_{e \in E} B_e) \leq \sum_{e \in E} \Pr(B_e). \tag{6}$$

If the number of edges are $< 2^{n-1}$ then we are sure that there is a 2-coloring of H where none of the events B_e occur. Clearly, this result took the extreme case. Already for $n = 2$ the result is not interesting, since it says that a graph with one vertex is 2-colorable. Let's apply the same argument to the case when an edge has at least k vertices. In this case the probability that an edge is monochromatic is $\leq 2^{1-k}$, and the number of edges are $\sum_{j=k}^{n} \binom{n}{j}$. Now if k is very small compared to n, then the number of edges is much larger than 2^{1-k} and the argument above fails. The reason is that the upper bound in (6) is not tight.

In PM we usually want to show that nothing "bad" occurs, i.e., if $B_1, \ldots, B_n$ were some bad events then we want to show that the probability that none of the bad events occur is positive, namely $\Pr(\overline{B}_1 \cap \overline{B}_2 \cap \cdots \cap \overline{B}_n) > 0$. Without making any assumptions on the dependence amongst the events, the best we can say about $\Pr(\overline{B}_1 \cap \overline{B}_2 \cap \cdots \cap \overline{B}_n)$ is that it is greater than $1 - \sum_i \Pr(B_i)$. However, this is meaningful if the sums of the probabilities is strictly smaller than one. On the other extreme, if we assume that the events $B_1, \ldots, B_n$ are **independent**, i.e., if

$$\Pr(B_1 \cap B_2 \cap \cdots \cap B_n) = \Pr(B_1)\Pr(B_2)\cdots\Pr(B_n),$$

then we can say something tighter, namely

$$\Pr(\overline{B}_1 \cap \overline{B}_2 \cap \cdots \cap \overline{B}_n) = \Pr(\overline{B}_1)\Pr(\overline{B}_2)\cdots\Pr(\overline{B}_n)$$

i.e., the events $\overline{B}_i$'s are also independent; note that the RHS is $> 1 - \sum_i \Pr(B_i)$. But what if there is some dependence amongst the events? What can we say then? How do we model these dependencies? Lovász Local Lemma (LLL) gives us a way to do that.

We first model the dependencies using directed graphs. Let $B_1, \ldots, B_n$ be events in a probability space. A directed graph $G = (V, E)$ with $V = [n]$ is a **dependency digraph** for $B_1, \ldots, B_n$ if each event B_i is independent of all the events B_j with $(i, j) \notin E$.

LEMMA 4 (Lovász Local Lemma (1975)). *Let $B_1, \ldots, B_n$ be events in a probability space and $G = (V, E)$ be their dependency digraph. Suppose there are n real numbers $x_1, \ldots, x_n$, $0 \leq x_i \leq 1$, s.t.*

$$\Pr(B_i) \leq x_i \prod_{(i,j) \in E} (1 - x_j) \tag{7}$$

then

$$\Pr(\bigcap_{i=1}^{n} \overline{B}_i) \geq \prod_{i=1}^{n} (1 - x_i).$$

Proof. The proof is based upon induction, but first we massage our claim. Repeatedly using the result that for two events A, B, $\Pr(A \cap B) = \Pr(B)\Pr(A|B)$, we obtain

$$\begin{aligned}
\Pr(\bigcap_{i=1}^{n} \overline{B}_i) &= \Pr(\overline{B}_1)\Pr(\bigcap_{i=2}^{n} \overline{B}_i|\overline{B}_1) \\
&= \Pr(\overline{B}_1)\Pr(\overline{B}_2|\overline{B}_1)\Pr(\bigcap_{i=2}^{n} \overline{B}_i|\overline{B}_1 \cap \overline{B_2}) \\
&= \Pr(\overline{B}_1)\Pr(\overline{B}_2|\overline{B}_1)\Pr(\overline{B}_3|\overline{B}_2 \cap \overline{B}_1)\Pr(\bigcap_{i=4}^{n} \overline{B}_i|\overline{B}_1 \cap \overline{B_2} \cap \overline{B_3}) \\
&\vdots \\
&= \Pr(\overline{B}_1)\Pr(\overline{B}_2|\overline{B}_1)\Pr(\overline{B}_3|\overline{B}_2 \cap \overline{B}_1)\ldots\Pr(\overline{B}_n|\bigcap_{i=1}^{n-1} \overline{B}_i).
\end{aligned}$$

We can further simplify the equation above by using the fact that $\Pr(A|B) + \Pr(\overline{A}|B) = 1$ (i.e., given B either A or its complement $\overline{A}$ must happen) to get

$$\Pr(\bigcap_{i=1}^{n} \overline{B}_i) = (1 - \Pr(B_1))(1 - \Pr(B_2|\overline{B}_1))(1 - \Pr(B_3|\overline{B}_2 \cap \overline{B}_1))\ldots(1 - \Pr(B_n|\bigcap_{i=1}^{n-1} \overline{B}_i)). \tag{8}$$

We next derive an upper bound on $\Pr(B_i| \cap_{j\in S} \overline{B}_j)$, for $S \subseteq [n]$. To get the desired result, we must show that

$$\Pr(B_i| \cap_{j\in S} \overline{B}_j) \leq x_i. \tag{9}$$

We prove this by using induction on $|S|$. When $|S| = 0$, then the claim follows from the assumption that $\Pr(B_i) \leq x_i \prod_{(i,j)\in E}(1 - x_j) \leq x_i$. So suppose the claim holds for all subsets of $[n]$ of size smaller than $|S|$. For succinctness, let $\overline{B}_S := \cap_{j\in S} \overline{B}_j$.

Consider $\Pr(B_i|\overline{B}_S)$. We can partition S into two parts S_1, S_2 where $S_1 = \{j \in S|(i, j) \in E\}$ and $S_2 = S \setminus S_1$. Then from the fact that $\Pr(A|B \cap C) = \Pr(A \cap B|C)/\Pr(B|C)$, we obtain

$$\Pr(B_i|\overline{B}_S) = \frac{\Pr(B_i \cap \overline{B}_{S_1}|\overline{B}_{S_2})}{\Pr(\overline{B}_{S_1}|\overline{B}_{S_2})}. \tag{10}$$

To get an upper bound on the LHS, we derive an upper bound on the numerator in the RHS and a lower bound on the denominator on the RHS.

1. Since $\Pr(A \cap B|C) \leq \Pr(A|C)$, we get

   $$\Pr(B_i \cap \overline{B}_{S_1}|\overline{B}_{S_2}) \leq \Pr(B_i|\overline{B}_{S_2}).$$

   But B_i is independent of B_{S_2} and hence also with its complement $\overline{B}_{S_2}$. Thus

   $$\Pr(B_i \cap \overline{B}_{S_1}|\overline{B}_{S_2}) \leq \Pr(B_i|\overline{B}_{S_2}) = \Pr(B_i) \leq x_i \prod_{(i,j)\in E}(1 - x_j), \tag{11}$$

   where the last step follows from (7). Note that if $S_1 = \emptyset$ then $\Pr(B_i|\overline{B}_S) = \Pr(B_i|\overline{B}_{S_2})$ and so this equation still applies, and we do not have to write it as a fraction as in (10). So in the next case, we can assume $|S_2| < |S|$.

2. Let $S_1 = \{j_1, \ldots, j_k\}$. Then from an argument similar to the one used to derive (8) but applied to conditional probabilities it follows that

   $$\Pr(\overline{B}_{S_1}|\overline{B}_{S_2}) = (1 - \Pr(B_{j_1}|\overline{B}_{S_2}))(1 - \Pr(B_{j_2}|\overline{B}_{j_1} \cap \overline{B}_{S_2}))\ldots(1 - \Pr(B_{j_k}|\overline{B}_{[k-1]} \cap \overline{B}_{S_2})).$$

   Since $|S_2| < |S|$, we apply (9) to each negative term on the RHS, to get

   $$\Pr(\overline{B}_{S_1}|\overline{B}_{S_2}) \geq \prod_{j\in S_1}(1 - x_j) \geq \prod_{j\in N(i)}(1 - x_j). \tag{12}$$

Plugging (11) and (12) in (10) we have the proof of the claim (9), hence completing the induction. Substituting the upper bound in (9) into (8) gives us the desired lower bound:

$$\Pr(\overline{B}_{[n]}) \geq \prod_{i=1}^{n}(1 - x_i).$$

Q.E.D.

As a useful corollary of LLL is the *symmetric-LLL*:

Corollary 5. *Let $B_1, \ldots, B_n$ be events s.t. $\Pr(B_i) \leq p$, for all $i = 1, \ldots, n$, and suppose that in the dependency digraph each event B_i has out-degree at most d, i.e., B_i is dependent on at most d other events. If $ep(d+1) \leq 1$ then with positive probability none of the events $B_1, \ldots, B_n$ occur.*

Proof. To get the claim we just have to choose x_i's such that $p \leq x_i \prod_{(i,j)\in E}(1-x_j)$. For simplicity, let's assume all x_i's are the same, say x. Suppose, for contradiction, $p > x \prod_{(i,j)\in E}(1-x)$. Since i has out-degree at most d, we have $p > x(1-x)^d$. Let's choose $x = 1/(d+1)$, then since $(1 - 1/(d+1))^d > 1/e$, we have $ep(d+1) > 1$, which is a contradiction to our assumption that $ep(d+1) \leq 1$. Thus $p \leq x \prod_{(i,j)\in E}(1-x)$, $x = 1/(d+1)$, and so with positive probability none of the events B_i's occur. **Q.E.D.**

Note: The corollary above is independent of the number of events. We will also need a simplified version for the asymmetric case:

Corollary 6. *Let $B_1, \ldots, B_n$ be events in a probability space and $G = (V, E)$ be their dependency graph. If for a given i, $\sum_{(i,j)\in E} \Pr(B_j) \leq 1/4$ then with positive probability none of the events $B_1, \ldots, B_n$ occur.*

Proof. Substitute $x_i := 2P(B_i)$. Then we want to show that $1 \leq 2\prod_{(i,j)\in E}(1 - 2\Pr(B_j))$. But we know that $\prod_{(i,j)\in E}(1 - 2\Pr(B_j)) \geq 1 - 2\sum_{(i,j)\in E}\Pr(B_j) \geq 1$ since $\sum_{(i,j)\in E}\Pr(B_j) \leq 1/4$. **Q.E.D.**

We now apply LLL to different problems.

¶1. Hypergraphs: What if instead of restricting edges, we restrict the dependency of each edge? The following theorem does that.

THEOREM 7. *Let H be a hypergraph in which every edges has at least k elements and each edge of H intersects at most d other edges. If $e(d+1) \leq 2^{k-1}$ then H is 2-colorable.*

Proof. The probability that an edge is monochromatic is $\leq 2^{1-k}$. Since each edge intersects at most d other edges, the event B_e is dependent on at most d other events. Since $e(d+1) \leq 2^{k-1}$, applying symmetric-LLL gives us the existence of a 2-colorable hypergraph. **Q.E.D.**

¶2. Directed Cycles of Given Multiplicity: Another surprising application of LLL is the following result:

THEOREM 8 (Alon and Linial (1989)). *Let $G = (V, E)$ be a simple directed graph with minimum out-degree δ and maximum indegree d. If $e(d\delta + 1)(1 - 1/k)^\delta < 1$ then D contains a directed simple cycle of length a multiple of k.*

Proof.

We will prove the theorem for the subgraph of G where the outdegree is *exactly* δ. Clearly, if the theorem holds for such a subgraph then it holds for G (WHY?).

To capture the periodicity of k, we use k-colors. The idea is to randomly k-color G and show that there is a cycle, where all the k colors appear in the cycle in a "consecutive manner" repeatedly.

Let $c : V \to \mathbb{Z}_k := \{0, \ldots, k-1\}$ be a random coloring of G, where each vertex $v \in V$ is colored uniformly at random with one of the k-colors $0, \ldots, k-1$. Let B_v be the even that there is no u, $(v, u) \in E$, s.t. $c(u) \equiv c(v) + 1 \mod k$.

1. $\Pr(A_v) \le (1-1/k)^\delta$, since the outdegree is at least δ and $(1-1/k)$ is the probability that a neighbor of v doesn't have the color $c(v)+1 \mod k$.

2. How many events A_u can be dependent with A_v? Two events A_v, A_u are dependent if the directed neighbourhood of v and u intersect. For each of the δ neighbors of v there are d vertices that point to it. Thus there are at most $d\delta$ events dependent with A_v.

3. By assumption, it follows that we can apply symmetric-LLL to get that there is a positive probability that none of the events happen, i.e., there is a coloring c such that for all $v \in V$ there is a u, $(v,u) \in E$, such that $c(u) \equiv c(v)+1 \mod k$.

4. Pick any vertex v_0 and form the sequence $v_0, v_1, v_2, \ldots$ such that

$$c(v_{j+1}) \equiv c(v_j) + 1 \mod k.$$

Suppose j is the smallest integer such that there exists an $\ell < j$ s.t. $v_\ell = v_j$. Then the cycle $v_\ell, \ldots, v_j$ is a simple cycle, and since the colors repeat at every kth. vertex, the length of the cycle is a multiple of k.

Q.E.D.

¶3. Frugal Colorings: A **proper coloring** of a graph G is a coloring of its vertices where neighbors get different colors. A proper coloring of G is said to be a **β-frugal coloring** if no color appears more than β times in the neighborhood of any vertex of G. When can we say that a graph has a β-frugal coloring?

We show the following result: If G has maximum degree $\Delta \ge \beta^\beta$ then it has a β-frugal coloring with roughly $c\Delta^{1+1/\beta}$ colors, for some constant $c \ge 2$.

For $\beta = 1$, the statement above says that we need $c\Delta^2$ colors to get a 1-frugal coloring, i.e., a coloring where every color appears at most once in the neighborhood of any vertex. One way to see this is to consider the square G^2 of G, i.e., the graph obtained by adding edges between vertices at distance at most two in G. Clearly, maximum degree of G^2 is smaller than Δ^2; moreover, every vertex in the neighborhood of G in G^2 is connected to every other vertex in the neighborhood. From Brook's theorem we know that a proper coloring of G^2 requires Δ^2+1 colors, and since $c \ge 2$, we have sufficient colors.

For $\beta \ge 2$, we pick a random (not necessarily proper) coloring of G with $c\Delta^{1+1/\beta}$ colors, where every color is picked with uniform probability $p := 1/c\Delta^{1+1/\beta}$. There are two types of events that prevent our coloring to be a proper β-frugal coloring:

Type-1. For each $(u,v) \in E$, let B_{uv} be the event that both u and v have the same color.

Type-2. For each set of $\beta+1$ neighbors $u_0, \ldots, u_\beta$ of some vertex, let $B_{u_0,\ldots,u_\beta}$ be the event that all these vertices have the same color.

What is the probability of these events taking place? What is the probability that some k vertices have the same color? It is p^{k-1}, since the probability is conditioned that one of the vertices gets a certain color, and the remaining get the same color.

Let's now look at the dependency digraph of these events. How many events is a type-1 event dependent upon? The vertices u,v are incident to at most 2Δ edges, and hence the same number of type-1 events. For type-2 events, suppose both u,v belong to the neighborhood $N(w)$ of a vertex w; then the type-2 event $B_{u_0,\ldots,u_\beta}$ where $\{u_0,\ldots,u_\beta\} \cap \{u,v\} \ne \emptyset$ is dependent on B_{uv}; there are $2\binom{\Delta}{\beta}$ such sets (there is some double counting here), and at most Δ choices for w (which happens when the neighborhood of both u and v are the same). How many events is a type-2 event dependent upon? Let $B_{u_0,\ldots,u_\beta}$ be a type-2 event. Then it is dependent on any type-1 event that has an endpoint in $u_0,\ldots,u_\beta$. Since maximum degree is Δ, there are $(\beta+1)\Delta$ type-1 events that $B_{u_0,\ldots,u_\beta}$ is dependent upon. The event $B_{u_0,\ldots,u_\beta}$ is dependent on another type-2 event $B_{v_0,\ldots,v_\beta}$ if $\{u_0,\ldots,u_\beta\} \cap \{v_0,\ldots,v_\beta\} \ne \emptyset$; for every u_i, $i = 0,\ldots,\beta$, there are Δ neighbors and for each neighbor there are $\binom{\Delta}{\beta}$ sets that contain u_i; thus $B_{u_0,\ldots,u_\beta}$ is dependent on $(\beta+1)\Delta\binom{\Delta}{\beta}$ type-2 events. Since the dependencies of type-2 events are larger, we will use those dependency bounds even for the type-1 events.

We now want to apply Corollary 6. In particular, we want to show that given some event B, the probability that one of the events dependent on B takes place is $\leq 1/4$. From the argument above, we know that B is dependent upon at most $(\beta+1)\Delta$ type-1 events and $(\beta+1)\Delta\binom{\Delta}{\beta}$ type-2 events. Therfore, the probability that some bad event in the neighborhood of B happens is bounded by

$$\begin{aligned}(\beta+1)\Delta p + (\beta+1)\Delta\binom{\Delta}{\beta}p^{\beta} &\leq (\beta+1)\Delta p + (\beta+1)\frac{\Delta^{\beta+1}}{\beta!}p^{\beta} \\ &\leq \frac{(\beta+1)\Delta}{c\Delta^{1+1/\beta}} + (\beta+1)\frac{\Delta^{\beta+1}}{\beta! c^{\beta}\Delta^{\beta+1}} \\ &\leq 1/4,\end{aligned}$$

where the last step follows if we choose $c \geq 16$.