Now we examine (similar to Theorem 1.4.2)

$$\Delta = \sum_{\alpha \sim \beta} \Pr[B_\alpha \wedge B_\beta]$$

We split the sum according to the number of *vertices* in the intersection of copies α and β. Suppose they intersect in j vertices. If $j = 0$ or $j = 1$ then $A_\alpha \cap A_\beta = \emptyset$ so that $\alpha \sim \beta$ cannot occur. For $2 \leq j \leq v$ let f_j be the maximal $|A_\alpha \cap A_\beta|$ where $\alpha \sim \beta$ and α, β intersect in j vertices. As $\alpha \neq \beta$, $f_v < e$. When $2 \leq j \leq v-1$ the critical observation is that $A_\alpha \cap A_\beta$ is a subgraph of H and hence, as H is strictly balanced,

$$\frac{f_j}{j} < \frac{e}{v}$$

There are $O(n^{2v-j})$ choices of α, β intersecting in j points since α, β are determined, except for order, by $2v - j$ points. For each such α, β

$$\Pr[B_\alpha \wedge B_\beta] = p^{|A_\alpha \cup A_\beta|} = p^{2e - |A_\alpha \cap A_\beta|} \leq p^{2e - f_j}$$

Thus

$$\Delta = \sum_{j=2}^{v} O(n^{2v-j}) O(n^{-\frac{v}{e}(2e-f_j)})$$

But

$$2v - j - \frac{v}{e}(2e - f_j) = \frac{v f_j}{e} - j < 0$$

so each term is $o(1)$ and hence $\Delta = o(1)$. By Janson's Inequality

$$\lim_{n \to \infty} \Pr[\wedge \overline{B}_\alpha] = \lim_{n \to \infty} M = exp[-c^e/a]$$

completing the proof. □

The fine threshold behavior for the appearance of an arbitrary graph H has been worked out but it can get quite complicated.

We set

$$p = \frac{\mu(1-\epsilon)}{\Delta}$$

so as to maximize this quantity. The added assumption of Theorem 1.2 assures us that the probability p is at most one. Then

$$E\left[-\ln[\Pr[\wedge_{i\in S}\overline{B_i}]\right] \geq \frac{\mu^2(1-\epsilon)}{2\Delta}$$

Therefore there is a specific $S \subset I$ for which

$$-\ln[\Pr[\wedge_{i\in S}\overline{B_i}] \geq \frac{\mu^2(1-\epsilon)}{2\Delta}$$

That is,

$$\Pr[\wedge_{i\in S}\overline{B_i}] \leq e^{-\frac{\mu^2(1-\epsilon)}{2\Delta}}$$

But

$$\Pr[\wedge_{i\in I}\overline{B_i}] \leq \Pr[\wedge_{i\in S}\overline{B_i}]$$

completing the proof. □

3 Appearance of Small Subgraphs Revisited

Generalizing the fine threshold behavior for the appearance of K_4 we find the fine threshold behavior for the appearance of any strictly balanced graph H.

Theorem 3.1 Let H be a *strictly* balanced graph with v vertices, e edges and a automorphisms. Let $c > 0$ be arbitrary. Let A be the property that G contains no copy of H. Then with $p = cn^{-v/e}$,

$$\lim_{n\to\infty} \Pr[G(n,p) \models A] = exp[-c^e/a]$$

Proof. Let $A_\alpha, 1 \leq \alpha \leq \binom{n}{v}v!/a$, range over the edge sets of possible copies of H and let B_α be the event $G(n,p) \supseteq A_\alpha$. We apply Janson's Inequality. As

$$\lim_{n\to\infty} \mu = lim_{n\to\infty}\binom{n}{v}v!p^e/a = c^e/a$$

we find

$$\lim_{n\to\infty} M = exp[-c^e/a]$$

Reversing

$$\begin{aligned} Pr[\overline{B_i}|\wedge_{1\le j<i}\overline{B_j}] &\le \Pr[\overline{B_i}]+\textstyle\sum_{j=1}^{d} Pr[B_j\wedge B_i] \\ &\le \Pr[\overline{B_i}]\left(1+\frac{1}{1-\epsilon}\sum_{j=1}^{d} Pr[B_j\wedge B_i]\right) \end{aligned}$$

since $\Pr[\overline{B_i}] \ge 1-\epsilon$. Employing the inequality $1+x \le e^x$,

$$Pr[\overline{B_i}|\wedge_{1\le j<i}\overline{B_j}] \le \Pr[\overline{B_i}]e^{\frac{1}{1-\epsilon}\sum_{j=1}^{d} Pr[B_j\wedge B_i]}$$

For each $1 \le i \le m$ we plug this inequality into

$$\Pr[\wedge_{i\in I}\overline{B_i}] = \prod_{i=1}^{m} Pr[\overline{B_i}|\wedge_{1\le j<i}\overline{B_j}]$$

The terms $\Pr[\overline{B_i}]$ multiply to M. The exponents add: for each $i,j \in I$ with $j<i$ and $j\sim i$ the term $Pr[B_j\wedge B_i]$ appears once so they add to $\Delta/2$. □

Proof of Theorem 1.2 As discussed earlier, the proof of Theorem 1.1 gives

$$\Pr[\wedge_{i\in I}\overline{B_i}] \le e^{-\mu+\frac{1}{1-\epsilon}\frac{\Delta}{2}}$$

which we rewrite as

$$-\ln[\Pr[\wedge_{i\in I}\overline{B_i}]] \ge \sum_{i\in I}\Pr[B_i] - \frac{1}{2(1-\epsilon)}\sum_{i\sim j}\Pr[B_i\wedge B_j]$$

For any set of indices $S\subset I$ the same inequality applied only to the B_i, $i\in S$ gives

$$-\ln[\Pr[\wedge_{i\in S}\overline{B_i}]] \ge \sum_{i\in S}\Pr[B_i] - \frac{1}{2(1-\epsilon)}\sum_{i,j\in S, i\sim j}\Pr[B_i\wedge B_j]$$

Let now S be a random subset of I given by

$$\Pr[i\in S] = p$$

with p a constant to be determined, the events mutually independent. (Here we are using probabilistic methods to prove a probability theorem!) Each term $\Pr[B_i]$ then appears with probability p and each term $\Pr[B_i\wedge B_j]$ with probability p^2 so that

$$\begin{aligned} &E\left[-\ln[\Pr[\wedge_{i\in S}\overline{B_i}]\right] \\ &\ge E\left[\textstyle\sum_{i\in S}\Pr[B_i]\right] - \frac{1}{2(1-\epsilon)}E\left[\sum_{i,j\in S,i\sim j}\Pr[B_i\wedge B_j]\right] \\ &= p\mu - \frac{1}{1-\epsilon}p^2\frac{\Delta}{2} \end{aligned}$$

2 The Proofs

The original proofs of Janson are based on estimates of the Laplace transform of an appropriate random variable. The proof presented here follows that of Boppana and Spencer [1989]. We shall use the inequalities

$$\Pr[B_i| \wedge_{j\in J} \overline{B_j}] \leq \Pr[B_i]$$

valid for all index sets $J \subset I, i \notin J$ and

$$\Pr[B_i|B_k \wedge \bigwedge_{j\in J} \overline{B_j}] \leq \Pr[B_i|B_k]$$

valid for all index sets $J \subset I, i, k \notin J$. The first follows from general Correlation Inequalities. The second is equivalent to the first since conditioning on B_k is the same as assuming $p_r = \Pr[r \in R] = 1$ for all $r \in A_k$.
Proof of Theorem 1.1 The lower bound follows immediately. Order the index set $I = \{1, \ldots, m\}$ for convenience. For $1 \leq i \leq m$

$$Pr[B_i| \wedge_{1\leq j<i} \overline{B_j}] \leq \Pr[B_i]$$

so

$$Pr[\overline{B_i}| \wedge_{1\leq j<i} \overline{B_j}] \geq Pr[\overline{B_i}]$$

and

$$\Pr[\wedge_{i\in I}\overline{B_i}] = \prod_{i=1}^{m} Pr[\overline{B_i}| \wedge_{1\leq j<i} \overline{B_j}] \geq \prod_{i=1}^{m} Pr[\overline{B_i}]$$

Now the upper bound. For a given i renumber, for convenience, so that $i \sim j$ for $1 \leq j \leq d$ and not for $d+1 \leq j < i$. We use the inequality $\Pr[A|B \wedge C] \geq \Pr[A \wedge B|C]$, valid for any A, B, C. With $A = B_i$, $B = \overline{B_1} \wedge \ldots \wedge \overline{B_d}$, $C = \overline{B_{d+1}} \wedge \ldots \wedge \overline{B_{i-1}}$

$$Pr[B_i| \wedge_{1\leq j<i} \overline{B_j}] = \Pr[A|B \wedge C] \geq Pr[A \wedge B|C] = Pr[A|C]Pr[B|A \wedge C]$$

From the mutual independence $Pr[A|C] = Pr[A]$. We bound

$$\Pr[B|A \wedge C] \geq 1 - \sum_{j=1}^{d} Pr[B_j|B_i \wedge C] \geq 1 - \sum_{j=1}^{d} Pr[B_j|B_i]$$

from the Correlation Inequality. Thus

$$Pr[B_i| \wedge_{1\leq j<i} \overline{B_j}] \geq \Pr[B_i] - \sum_{j=1}^{d} Pr[B_j \wedge B_i]$$

the value of $\Pr[\wedge_{i\in I}\overline{B_i}]$ if the B_i were independent.
Theorem 1.1 (The Janson Inequality). Let $B_i, i \in I$, Δ, M be as above and assume all $\Pr[B_i] \leq \epsilon$. Then

$$M \leq \Pr[\wedge_{i\in I}\overline{B_i}] \leq Me^{\frac{1}{1-\epsilon}\frac{\Delta}{2}}$$

Now set

$$\mu = E[X] = \sum_{i\in I} \Pr[B_i]$$

For each $i \in I$

$$\Pr[\overline{B_i}] = 1 - \Pr[B_i] \leq e^{-\Pr[B_i]}$$

so, multiplying over $i \in I$,

$$M \leq e^{-\mu}$$

It is often more convenient to replace the upper bound of Theorem 1.1 with

$$\Pr[\wedge_{i\in I}\overline{B_i}] \leq e^{-\mu+\frac{1}{1-\epsilon}\frac{\Delta}{2}}$$

As an example, set $p = cn^{-2/3}$ and consider the probability that $G(n,p)$ contains no K_4. The B_i then range over the $\binom{n}{4}$ potential K_4 - each being a 6-element subset of Ω. Here, as is often the case, $\epsilon = o(1)$, $\Delta = o(1)$ (as calculated previously) and μ approaches a constant, here $k = c^6/24$. In these instances $\Pr[\wedge_{i\in I}\overline{B_i}] \to e^{-k}$. Thus we have the fine structure of the threshold function of $\omega(G) = 4$.

As Δ becomes large the Janson Inequality becomes less precise. Indeed, when $\Delta \geq 2\mu(1-\epsilon)$ it gives an upper bound for the probability which is larger than one. At that point (and even somewhat before) the following result kicks in.
Theorem 1.2 (The Generalized Janson Inequality). Under the assumptions of Theorem 1.1 and the further assumption that $\Delta \geq \mu(1-\epsilon)$

$$\Pr[\wedge_{i\in I}\overline{B_i}] \leq e^{-\frac{\mu^2(1-\epsilon)}{2\Delta}}$$

Theorem 1.2 (when it applies) often gives a much stronger result than Chebyschev's Inequality as used earlier. We can bound $Var[X] \leq \mu + \Delta$ so that

$$\Pr[\wedge_{i\in I}\overline{B_i}] = \Pr[X = 0] \leq \frac{Var[X]}{E[X]^2} \leq \frac{\mu+\Delta}{\mu^2}$$

Suppose $\epsilon = o(1)$, $\mu \to \infty$, $\mu \ll \Delta$, and $\gamma = \frac{\mu^2}{\Delta} \to \infty$. Chebyschev's upper bound on $\Pr[X = 0]$ is then roughly γ^{-1} while Janson's upper bound is roughly $e^{-\gamma}$.

Lecture 3: The Poisson Paradigm

When X is the sum of many rare indicator "mostly independent" random variables and $\mu = E[X]$ we would like to say that X is close to a Poisson distribution with mean μ and, in particular, that $\Pr[X = 0]$ is nearly $e^{-\mu}$. We call this rough statement the Poisson Paradigm. We give a number of situations in which this Paradigm may be rigorously proven.

1 The Janson Inequalities

In many instances we would like to bound the probability that none of a set of bad events $B_i, i \in I$ occur. If the events are mutually independent then

$$\Pr[\wedge_{i\in I}\overline{B_i}] = \prod_{i\in I}\Pr[\overline{B_i}]$$

When the B_i are "mostly" independent the Janson Inequalities allow us, sometimes, to say that these two quantities are "nearly" equal.

Let Ω be a finite universal set and let R be a random subset of Ω given by

$$\Pr[r \in R] = p_r,$$

these events mutually independent over $r \in \Omega$. (In application to $G(n,p)$, Ω is the set of pairs $\{i,j\}$, $i,j \in V(G)$ and all $p_r = p$ so that R is the edge set of $G(n,p)$.) Let $A_i, i \in I$, be subsets of Ω, I a finite index set. Let B_i be the *event* $A_i \subseteq R$. (That is, each point $r \in \Omega$ " flips a coin" to determine if it is in R. B_i is the event that the coins for all $r \in A_i$ came up "heads".) Let X_i be the indicator random variable for B_i and $X = \sum_{i\in I} X_i$ the number of $A_i \subseteq R$. The event $\wedge_{i\in I}\overline{B_i}$ and $X = 0$ are then identical. For $i,j \in I$ we write $i \sim j$ if $i \neq j$ and $A_i \cap A_j \neq \emptyset$. Note that when $i \neq j$ and not $i \sim j$ then B_i, B_j are independent events since they involve separate coin flips. Furthermore, and this plays a crucial role in the proofs, if $i \notin J \subset I$ and not $i \sim j$ for all $j \in J$ then B_i is mutually independent of $\{B_j | j \in J\}$, i.e., independent of any Boolean function of those B_j. This is because the coin flips on A_i and on $\cup_{j\in J} A_j$ are independent. We define

$$\Delta = \sum_{i\sim j}\Pr[B_i \wedge B_j]$$

Here the sum is over ordered pairs so that $\Delta/2$ gives the same sum over unordered pairs. (This will be the same Δ as in Lecture 1. We set

$$M = \prod_{i\in I}\Pr[\overline{B_i}],$$