

Random Sums and Graphs

Lecture 1

Probability

Events, probabilities,
conditional probabilities,
union bound

Random variables,
expected value,
linearity of expectation
Markov's inequality

Algorithms

Approx. 3-SAT,

Big idea: we only want the average

$$\begin{array}{l} \text{Linearity of Expectation} \\ E[X+Y] = E[X] + E[Y] \quad (\text{stochastic variable}) \\ \text{"average sum = sum of averages"} \end{array}$$

let $Y_i = \begin{cases} 1 & \text{if } i\text{th clause satisfied} \\ 0 & \text{otherwise} \end{cases} \quad (i=1, \dots, m)$

$$E[Y_i] = 7/8 \quad (1 \text{ clause case})$$

$$E[\text{\# clauses satisfied}] = E\left[\sum_{i=1}^m Y_i\right] = \sum_{i=1}^m E[Y_i] = 7/8 m$$

let $Y_i = \begin{cases} 1 & \text{if } i\text{th clause satisfied} \\ 0 & \text{otherwise} \end{cases}$

$$Z = \sum_{i=1}^m Y_i = \text{\# clauses satisfied}$$

$$E[Z] = \frac{1}{2} E[Z | x_1 = \text{true}] + \frac{1}{2} E[Z | x_1 = \text{false}]$$

$\Rightarrow$ either $E[Z | x_1 = \text{true}]$ or $E[Z | x_1 = \text{false}]$ is $\geq E[Z]$!

QuickSort,

Random divide and conquer

① pick random 'pivot'

② Divide A into two parts



Expected running time of QuickSort

For each $i, j \in [n]$ w/ $i < j$,

$X_{ij} = \begin{cases} 1 & \text{if } i\text{th smallest element is compared to } j\text{th smallest} \\ 0 & \text{otherwise} \end{cases}$

$\sum_{i,j} X_{ij}$ = total # comparisons made by the algorithm
 $\approx$ total running time

$E[\sum_{i,j} X_{ij}]$ = average running time

$$E[\sum_{i,j} X_{ij}] = \sum_{i,j} E[X_{ij}]$$

now analyze $E[X_{ij}]$ in isolation

$$E[X_{ij}] = \Pr[X_{ij} = 1] = \Pr[i\text{th element is compared to } j\text{th}]$$

i th and j th elem selected before any element in between happens w/ probability $\frac{2}{j-i+1}$

$$E[X_{ij}] = \frac{2}{j-i+1}$$

$$E[\sum_{i,j} X_{ij}] = \sum_{i,j} E[X_{ij}] = \sum_{i,j} \frac{2}{j-i+1}$$

$$\leq 2n \left(\frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{n} \right) \leq O(n \log n) \quad \text{with harmonic series}$$

Quickselect (HW)

Lecture 2

Probability

Hash Functions,

Universal Hash Functions

$h: [n] \rightarrow k$ s.t. for all pairs $i, j \in [n]$,
 $\Pr[h(i) = h(j)] \leq \frac{1}{k}$ (much weaker than ideal hash)

e.g. $h(x) = (ax + b) \bmod p \bmod k$

where • p prime, $> k$

• a in $\{1, \dots, p-1\}$ unif. random

• b in $\{0, \dots, p-1\}$ unif. random



Markov's Inequality



Markov's Inequality (special case)

Let $X \geq 0$ be nonnegative RV

$$\Pr[X \geq 2E[X]] \leq \frac{1}{2}$$

Algorithms

Frequency Estimation, Heavy Hitters (in streams)



estimate f_i w/
 $\min_{j=1, \dots, d} A_j[h_j(i)]$

Lecture 3 (last Tuesday)

Probability



k-wise independence,
higher order moments
(generalizing universal hash
functions and Markov's inequality)

Algorithms

Hash tables w/
chaining +
linear probing

k-wise independent hash functions
 $h: [U] \rightarrow [m]$
s.t. for any k keys $x_1, x_2, \dots, x_k \in [U]$,
 k values $v_1, v_2, \dots, v_k \in [m]$,
 $P[h(x_1)=v_1, h(x_2)=v_2, \dots, h(x_k)=v_k] = \frac{1}{m^k}$
eg. $h(x) = a_0 + a_1x + a_2x^2 + \dots + a_{k-1}x^{k-1} \pmod p$
where $p \geq U$ prime,
 $a_0, a_1, \dots, a_{k-1} \in [p]$ indep, uniformly random


Important lemma
k-wise independent random variables
 $P[X_1=y_1, X_2=y_2, \dots, X_k=y_k]$
 $= P[X_1=y_1] P[X_2=y_2] \dots P[X_k=y_k]$
for all values $y_1, y_2, \dots, y_k$
let $X_1, \dots, X_n \in \{0,1\}$ be 4-wise independent
and $\mu = E[X_1 + \dots + X_n]$. then for $\beta > 0$,
 $P[X_1 + \dots + X_n \geq \mu + \beta] \leq \frac{4 \max\{\mu, \mu^3\}}{\beta^4}$



Lecture 4 (last Thursday)


Probability

Galton-Watson processes



Algorithms


Random Contractions for min cut



Randomized contractions

Input: $G=(V,E)$, weights $c \in \mathbb{R}_{>0}^E$ (unweighted)

- ① while $|E| > 1$
 - Ⓐ sample $e \in E$ in proportion to c
 - Ⓑ "contract" e
- ② return set of edges (in original graph) contracted to remaining edge




Amplification by branching

Instead of restarting at the beginning, restart partway through

High-level algorithm

- ① randomly contract until prob. of ruining the min-cut is $\sim 1/2$
- ② Branch: make two copies of current graph recurse on both copies



Lecture 5 (today)

Probability

Multiplicative
Chernoff Bound

Algorithms (/ results)

Random graphs
& the giant component

Lecture 5 (today)

Probability

Multiplicative
Chernoff Bound

Algorithms (/ results)

Random graphs
& the giant component



Flip 100 coins.

50 come up heads.

Should you be surprised?

(Maybe a little: $\approx 8\%$ chance)



Flip 100 coins.

0 come up heads.

Should you be surprised?

Absolutely: $\frac{1}{2^{100}}$ chance



Flip 100 coins.

≤ 25 come up heads.

Should you be surprised?

Yes: $\approx 2.818 \times 10^{-7}$



Flip 10 coins.

≤ 4 come up heads.

Should you be surprised?

No: $\approx 37.7\%$



Flip 1000 coins.

$\leq 40\%$ come up heads.

Should you be surprised?

Yes: $\approx 1.364 \times 10^{-10}$ probability



10 coins.

≤ 4 heads.

$\approx 37.7\%$

1000 coins.

$\leq 40\%$ heads.

1.364×10^{-10} prob.

scale matters

Concentration phenomena

let $X_1, \dots, X_n \in [0, 1]$ be independent

$$\mu = E[X_1 + \dots + X_n]$$

Q: probability that

$X_1 + \dots + X_n$ is close to μ ?
(multiplicatively)

A: error probability exponentially small in n

Multiplicative Chernoff bound

let $X_1, \dots, X_n \in [0, 1]$ be independent

$$\mu = E[X_1 + \dots + X_n]$$

$$\epsilon \in (0, 1)$$

$$\bullet \quad P[X_1 + \dots + X_n \geq (1 + \epsilon)\mu] \leq e^{-\epsilon^2 \mu / 3}$$

$$\bullet \quad P[X_1 + \dots + X_n \leq (1 - \epsilon)\mu] \leq e^{-\epsilon^2 \mu / 2}$$

(proof at end if time permits)

Lecture 4 (last Thursday)

Probability

Multiplicative
Chernoff Bound

Algorithms (/ results)

Random graphs
& the giant component


Erdős-Rényi graphs $G(n, p)$

(1959)

- n vertices
- each edge sampled independently
w/ probability p


Q: what does such a graph
look like?

does it consistently have
certain properties?



Erdős-Rényi graphs $G(n, p)$

- n vertices
- each edge sampled independently w/ probability p



Theorem 5.3. Consider a random graph $G \sim \mathcal{G}(n, p)$ for $p = c/n$, where c is a constant.

1. If $c > 1$, then with high probability, there is exactly one connected component of G with $\Omega(n)$ vertices, and all other components have size $\leq O(\log n)$.
2. For $c < 1$, then with high probability, all connected components of G has size $< O(\log n)$.

"threshold phenomena"

Theorem 5.3. Consider a random graph $G \sim \mathcal{G}(n, p)$ for $p = c/n$, where c is a constant.

1. If $c > 1$, then with high probability, there is exactly one connected component of G with $\Omega(n)$ vertices, and all other components have size $\leq O(\log n)$.
2. For $c < 1$, then with high probability, all connected components of G has size $< O(\log n)$.

Proof outline ($c > 1$)

- ① Gap theorem let $p = (1 + \epsilon)/n$. w/ high probability, all components have size $\leq O\left(\frac{\log n}{\epsilon^2}\right)$ or $\geq \Omega(\epsilon n)$
- ② Existence of large component
- ③ Uniqueness of large component


① Gap theorem let $p = (1 + \epsilon)/n$. w/ high probability,
all components have size

$$\leq O\left(\frac{\log n}{\epsilon^2}\right) \quad \text{or} \quad \geq \Omega(\epsilon n)$$

Setup

fix vertex v .

let $C(v) \subseteq V$ = connected component of v



to analyze $C(v)$, imagine
revealing $C(v)$ w/ a search algo

Fix vertex v . $C(v) \subseteq V =$ connected component of v

"Search algo"


$A = \{\text{vertices known to be in } C(v)\}$ (initially $A = \{v\}$)

$B = \{\text{explored vertices}\}$ (initially $B = \emptyset$)

① while there is a vertex $v_i \in A \setminus B$

 ② add all neighbors of v_i to A

 ③ add v_i to B



Indexed version of "Search algo"

$A_i = \{\text{vertices known to be in } C(v) \text{ after } i \text{ iter.}\}$
($A_0 = \{v\}$)

$B_i = \{\text{explored vertices after } i \text{ iterations}\}$ ($B_0 = \emptyset$)

1. For $i = 1, 2, \dots$ until $A_i = B_i$

a. Let $v_i \in A_{i-1} \setminus B_{i-1}$

b. $A_i = A_{i-1} \cup N(v_i)$

c. $B_i = B_{i-1} \cup \{v_i\}$

Observation 1

algo terminates iff $|A_i| = i$

$A_i = \{\text{vertices known to be in } C(v) \text{ after } i \text{ iter.}\}$
($A_0 = \{v\}$)


$B_i = \{\text{explored vertices after } i \text{ iterations}\}$ ($B_0 = \emptyset$)

1. For $i=1, 2, \dots$ until $A_i = B_i$

a. Let $v_i \in A_{i-1} \setminus B_{i-1}$

b. $A_i = A_{i-1} \cup N(v_i)$

c. $B_i = B_{i-1} \cup \{v_i\}$



Observation 1

algo terminates iff $|A_i| = i$

Observation 2 $A_{i+1} = A_i \cup S_{i+1}$

where S_{i+1} samples each $v \in V \setminus A_i$ independently
w/ probability p

$A_i = \{\text{vertices known to be in } C(v) \text{ after } i \text{ iter.}\}$
($A_0 = \{v\}$)

$B_i = \{\text{explored vertices after } i \text{ iterations}\}$ ($B_0 = \emptyset$)

1. For $i=1, 2, \dots$ until $A_i = B_i$

a. Let $v_i \in A_{i-1} \setminus B_{i-1}$

b. $A_i = A_{i-1} \cup N(v_i)$

c. $B_i = B_{i-1} \cup \{v_i\}$

Observation 1 algo terminates iff $|A_i| = i$

Observation 2 $A_{i+1} = A_i \cup S_{i+1}$

where S_{i+1} samples each $v \in V \setminus A_i$ independently w/ probability p

1. For $i=1,2,\dots$ until $A_i = B_i$

a. Let $v_i \in A_{i-1} \setminus B_{i-1}$

b. $A_i = A_{i-1} \cup N(v_i)$

c. $B_i = B_{i-1} \cup \{v_i\}$

1. For $i=1,2,\dots$ until ~~$|A_i| = i$~~

a. Let S_i randomly sample $V \setminus A_{i-1}$ w/ probability p

b. $A_i = A_{i-1} \cup S_i$

(equivalent)

$$e^x = 1 + x + \frac{x^2}{2} + \frac{x^3}{3!} + \frac{x^4}{4!} + \dots$$

let $i \leq \frac{\epsilon n}{4}$.

Lemma 1. $E[|A_i|] \geq (1+\epsilon)i$

$E[|A_0|] = 1$

$E[|A_1|] = 1 + (n-1)p$

$E[|A_2|] = 1 +$

$P_r[u \notin A_1] = (1-p)$

$P_r[u \notin A_2] = (1-p)^2$

$E[|A_i|] - 1 = (n-1)(1 - (1-p)^i)$

$\geq (n-1) \cdot (1 - \frac{\epsilon}{4}) \cdot ip \geq (1+\epsilon)i$

① Gap theorem let $p = (1+\epsilon)/n$. w/ high probability, all components have size $\leq O(\frac{\log n}{\epsilon^2})$ or $\geq \Omega(\epsilon n)$

Multiplicative Chernoff bound
let $X_1, \dots, X_n \in [0,1]$ be independent

$\mu = E[X_1 + \dots + X_n] \in (0,1)$

$P[X_1 + \dots + X_n \leq (1-\epsilon)\mu] \leq e^{-\epsilon^2 \mu / 2}$

1. For $i=1, 2, \dots$ until $|A_i| \geq i$

a. Let S_i randomly sample $V \setminus A_{i-1}$ w/ probability p

b. $A_i = A_{i-1} \cup S_i$

$(1-p)^i \leq e^{-pi}$ $ip \leq \frac{\epsilon n (1+\epsilon)}{4n}$

$\leq 1 - ip + (ip)^2$

$\leq 1 - ip + \frac{\epsilon ip}{2} = 1 - (1 - \frac{\epsilon}{4})ip$

$ip \geq \frac{1+\epsilon}{4}$

let $i \leq \frac{\epsilon n}{4}$.

Lemma 1. $E[|A_i|] \geq (1+\epsilon)i$

Lemma 2.

$$P[|A_i| \leq i] \leq e^{-\epsilon^2 i / 8}$$

$$\approx (1-\epsilon)E[|A_i|]$$

$$X_i = \begin{cases} 1 & \text{if } i\text{th vertex is in } A_i \\ 0 & \text{otherwise} \end{cases}$$

$$\leq P[X_1 + \dots + X_n \leq (1-\frac{\epsilon}{2})E[|A_i|]] \leq e^{-\frac{(\epsilon/2)^2 i}{2}} = e^{-\epsilon^2 i / 8}$$

① Gap theorem let $p = (1+\epsilon)/n$. w/ high probability, all components have size $\leq O(\frac{\log n}{\epsilon^2})$ or $\geq \Omega(\epsilon n)$

Multiplicative Chernoff bound
let $X_1, \dots, X_n \in [0,1]$ be independent

$$\mu = E[X_1 + \dots + X_n] \in (0,1)$$

$$P[X_1 + \dots + X_n \leq (1-\epsilon)\mu] \leq e^{-\epsilon^2 \mu / 2}$$

1. For $i=1, 2, \dots$ until $|A_i| \geq i$

a. Let S_i randomly sample $V \setminus A_{i-1}$ w/ probability p

$$b. A_i = A_{i-1} \cup S_i$$

let $i \leq \frac{\epsilon n}{4}$.

Lemma 1. $E[|A_i|] \geq (1+\epsilon)i$

Lemma 2.

$$P[|A_i| \leq i] \leq e^{-\epsilon^2 i / 8}$$

for $i \geq 32 \log n / \epsilon^2$, $\uparrow \leq \boxed{\frac{1}{n^4}}$

by union bound over $\frac{32 \log n}{\epsilon^2} \leq i \leq \frac{\epsilon n}{4}$,

$$P[|A_i| \leq i \text{ for any } i \text{ s.t. } \frac{32 \log n}{\epsilon^2} \leq i \leq \frac{\epsilon n}{4}] \leq \frac{1}{n^3}$$

① Gap theorem let $p = (1+\epsilon)/n$. w/ high probability, all components have size $\leq O(\frac{\log n}{\epsilon^2})$ or $\geq \Omega(\epsilon n)$

Multiplicative Chernoff bound
let $X_1, \dots, X_n \in [0, 1]$ be independent

$$\mu = E[X_1 + \dots + X_n] \quad \epsilon \in (0, 1)$$

$$P[X_1 + \dots + X_n \leq (1-\epsilon)\mu] \leq e^{-\epsilon^2 \mu / 2}$$

1. For $i=1, 2, \dots$ until $|A_i|=i$
 - a. Let S_i randomly sample $V \setminus A_{i-1}$ w/ probability p
 - b. $A_i = A_{i-1} \cup S_i$

Theorem 5.3. Consider a random graph $G \sim \mathcal{G}(n, p)$ for $p = c/n$, where c is a constant.

1. If $c > 1$, then with high probability, there is exactly one connected component of G with $\Omega(n)$ vertices, and all other components have size $\leq O(\log n)$.
2. For $c < 1$, then with high probability, all connected components of G has size $< O(\log n)$.

Proof outline ($c > 1$)

✓ ① Gap theorem let $p = (1 + \epsilon)/n$. w/ high probability,
all components have size

$$\leq O\left(\frac{\log n}{\epsilon^2}\right) \quad \text{or} \quad \geq \Omega(\epsilon n)$$

② Existence of large component


③ Uniqueness of large component

② Existence of large component

Lemma 5.5. *Let $p = (1 + \epsilon)/n$ for $\epsilon > 0$. Let $v \in V$ be a vertex. For all $3 \leq h \leq \epsilon n$, with probability at least $1/h$, v has at least $1 + h$ vertices in its connected component.*

Lemma 5.5. Let $p = (1 + \epsilon)/n$ for $\epsilon > 0$. Let $v \in V$ be a vertex. For all $3 \leq h \leq \epsilon n$, with probability at least $1/h$, v has at least $1 + h$ vertices in its connected component.

Last time




let T = complete binary tree of height h ,
w/ every edge deleted w/ prob $\leq 1/2$

Fact: $P[\text{some leaf connected to root}] \geq \frac{1}{h}$

"Galton-Watson process"

Lemma 5.5. Let $p = (1 + \epsilon)/n$ for $\epsilon > 0$. Let $v \in V$ be a vertex. For all $3 \leq h \leq \epsilon n$, with probability at least $1/h$, v has at least $1 + h$ vertices in its connected component.




let $T =$ complete k -ary tree of height h ,
 w/ every edge ^{survives} ~~deleted~~ w/ prob $\geq 1/k$

Fact: $P[\text{some leaf connected to root}] \geq \frac{1}{h}$

"Galton-Watson process"

Lemma 5.5. Let $p = (1 + \epsilon)/n$ for $\epsilon > 0$. Let $v \in V$ be a vertex. For all $3 \leq h \leq \epsilon n$, with probability at least $1/h$, v has at least $1 + h$ vertices in its connected component.



let T = complete k -ary tree of height h ,
w/ every edge deleted w/ prob $\leq 1/k$

Fact: $P[\text{some leaf connected to root}] \geq \frac{1}{h}$

"Galton-Watson process"

Equivalent process

start w/ population of size 1

Each generation, each individual flips k coins
w/ prob. $1/k$ of heads

total # heads gives size of next generation


Fact: $P[\text{surviving } h \text{ generations}] \geq \frac{1}{h}$

Lemma 5.5. Let $p = (1 + \epsilon)/n$ for $\epsilon > 0$. Let $v \in V$ be a vertex. For all $3 \leq h \leq \epsilon n$, with probability at least $1/h$, v has at least $1 + h$ vertices in its connected component.

Recall search algo.

conditional on $|A_{i-1}| \leq h$,

A_i/A_{i-1} is like sampling
 $n-h \geq \frac{n}{1+\epsilon}$ children each
w/ prob. $p = \frac{1+\epsilon}{n}$



let T = complete k -ary tree of height h ,
w/ every edge deleted w/ prob $\leq 1/k$

Fact: $P[\text{some leaf connected to root}] \geq \frac{1}{h}$
"Galton-Watson process"

1. For $i=1, 2, \dots$ until $A_i = B_i$
 - a. Let $v_i \in A_{i-1} \setminus B_{i-1}$
 - b. $A_i = A_{i-1} \cup N(v_i)$
 - c. $B_i = B_{i-1} \cup \{v_i\}$

surviving h generations $\Rightarrow |A_i| > i \quad \forall i < h$
(or reaching $|A_i| = h$)

Theorem 5.3. Consider a random graph $G \sim \mathcal{G}(n, p)$ for $p = c/n$, where c is a constant.

Proof outline ($c > 1$)

1. If $c > 1$, then with high probability, there is exactly one connected component of G with $\Omega(n)$ vertices, and all other components have size $\leq O(\log n)$.
2. For $c < 1$, then with high probability, all connected components of G has size $< O(\log n)$.

✓ ① Gap theorem let $p = (1 + \varepsilon)/n$. w/ high probability,
all components have size
 $\leq O\left(\frac{\log n}{\varepsilon^2}\right)$ or $\geq \Omega(\varepsilon n)$

✓ ② Existence of large component


③ Uniqueness of large component

Imagine sampling $G \sim (n, p)$ as:

- ① sample $G_1 \sim (n/2, p)$
- ② sample $G_2 \sim (n/2, p)$
- ③ combine G_1 and G_2
- ④ sample cross edges

$$\frac{\varepsilon^2 n^2}{4} \cdot p = \underline{\varepsilon^2 n}$$

($\leq \frac{1}{\varepsilon}$ large
comp,
 ≥ 1)



Multiplicative Chernoff bound

let $X_1, \dots, X_n \in [0, 1]$ be independent


$$\mu = E[X_1 + \dots + X_n]$$

$$\epsilon \in (0, 1)$$

$$\bullet \quad P[X_1 + \dots + X_n \geq (1 + \epsilon)\mu] \leq e^{-\epsilon^2 \mu / 3}$$

$$\bullet \quad P[X_1 + \dots + X_n \leq (1 - \epsilon)\mu] \leq e^{-\epsilon^2 \mu / 2}$$

(proof at end if time permits)



Chapter 5

Random Sums and Graphs


5.1 Random sums

If, out of 100 coin tosses, you were told that 50 of them were heads, would you be surprised? Actually, you should be a little surprised. The odds of getting exactly 50 heads is about 8%. But if you were told that the number was in the range 45 to 55, you probably wouldn't think much of it.

If you were told that all 100 coin tosses came up heads, you wouldn't believe it. The odds of that, we know, is $1/2^{100}$. If you bet money and lost on this event, you would be outraged (and, at even odds, certainly broke for the rest of eternity).

Suppose you were told that at most 25 coin tosses came up heads. Should you be surprised? On one hand, 25 is half of the expected amount. On the other hand, the claim is not that there was exactly 25 heads, but *at most* 25 heads. There could be 25, 24, 23, etc., down to 0. Even though the event of getting any one of these counts should be low, being far from average, there are also 26 of these events. Do the probabilities add up to very much? It turns out that the probability of getting 25 or fewer heads is tiny: about 2.818×10^{-7} .

The *scale* is very important in this discussion. If, out of 10 coin tosses, you got 4 or fewer heads, you shouldn't be too surprised. There is (roughly) a 37.7% chance of getting at most 4 heads. But if at most 40% of 1000 coin tosses came up heads, you should be *very* surprised. The odds of this occurring is occurring is roughly 1.364×10^{-10} .



We generalize the discussion to coins with any fixed probability of heads, $p \in [0, 1]$. The *binomial distribution*, denoted $\mathcal{B}(n, p)$, is the distribution of the number of heads over n independent coin tosses that each flip heads with probability p . The probabilities of different binomial distributions is plotted above. (See also [Wikib](#).)

We write $B \sim \mathcal{B}(n, p)$ to denote a random variable $B \in \{0, \dots, n\}$ drawn from the binomial distribution $\mathcal{B}(n, p)$. The expected value of B is $\mathbf{E}[B] = pn$. The following lemma bounds the probability of B being a multiplicative factor smaller than its mean, pn . Note that the probability decays *exponential fast in the mean*.

Lemma 5.1. *Let $B \sim \mathcal{B}(n, p)$ and $\epsilon \in (0, 1)$. Then*

$$\mathbf{P}[B \leq (1 - \epsilon)pn] \leq e^{-\epsilon^2 pn/2}.$$

Proof. We have

$$\mathbf{P}[B \leq (1 - \epsilon)pn] = \mathbf{P}[e^{-\epsilon B} \geq e^{-\epsilon(1-\epsilon)pn}] \stackrel{(a)}{\leq} \frac{\mathbf{E}[e^{-\epsilon B}]}{e^{-\epsilon(1-\epsilon)pn}} = e^{\epsilon(1-\epsilon)pn} \mathbf{E}[e^{-\epsilon B}] \quad (5.1)$$

by (a) Markov's inequality. It remains to analyze $\mathbf{E}[e^{-\epsilon B}]$. Write B as the sum $B = X_1 + \dots + X_n$, where each X_i is an independent $\{0, 1\}$ -random variable with $\mathbf{P}[X_i = 1] = p$. Now we have

$$\mathbf{E}[e^{-\epsilon B}] = \mathbf{E}[e^{-\epsilon X_1 - \epsilon X_2 - \dots - \epsilon X_n}] \stackrel{(b)}{=} \mathbf{E}[e^{-\epsilon X_1}] \mathbf{E}[e^{-\epsilon X_2}] \dots \mathbf{E}[e^{-\epsilon X_n}],$$

where (b) is by independent of the X_i 's. For each X_i , we have

$$\begin{aligned} \mathbf{E}[e^{-\epsilon X_i}] &= pe^{-\epsilon} + (1 - p) = 1 + p(e^{-\epsilon} - 1) \\ &\stackrel{(c)}{\leq} p(1 - \epsilon + \epsilon^2/2) + (1 - p) = 1 - (\epsilon - \epsilon^2/2)p \\ &\stackrel{(d)}{\leq} e^{-(\epsilon - \epsilon^2/2)p}. \end{aligned}$$

Here (c) uses the inequality $e^{-x} \leq 1 - x + x^2/2$ for all $x > 0$ ¹. (d) is by the inequality $1 + x \leq e^x$ for all x . Thus,

$$\mathbf{E}[e^{-\epsilon B}] = \mathbf{E}[e^{-\epsilon X_1}] \mathbf{E}[e^{-\epsilon X_2}] \cdots \mathbf{E}[e^{-\epsilon X_n}] \leq e^{-(\epsilon - \epsilon^2/2)pn}, \quad (5.2)$$

Putting everything together, we have

$$\mathbf{P}[B \leq (1 - \epsilon)pn] \stackrel{(e)}{\leq} e^{\epsilon(1-\epsilon)pn} \mathbf{E}[e^{-\epsilon B}] \stackrel{(f)}{\leq} e^{\epsilon(1-\epsilon)pn - (\epsilon - \epsilon^2/2)pn} = e^{-\epsilon^2 pn/2},$$

by (e) inequality (5.1) and (f) inequality (5.2), as desired. $\square$

One can prove a similar inequality bounding the probability that B exceeds its mean by a multiplicative factor. The proof is similar to lemma 5.1 and left as exercise 5.1.

Lemma 5.2. *Let $B \sim \mathcal{B}(n, p)$ and $\epsilon \in (0, 1)$. Then*

$$\mathbf{P}[B \geq (1 + \epsilon)pn] \leq e^{-\epsilon^2 pn/3}.$$

5.2 Random graphs

Paul Erdős, inspired by Ramsey [Ram30] before him, had a series of work analyzing *random graphs*, producing a large body of results that can mostly be grouped into

¹To see that

$$1 - x + x^2/2 \geq e^{-x}$$

for all $x \geq 0$, observe first that both sides equal 1 at $x = 0$. The derivative of the LHS, $-(1 - x)$, is always at least the RHS of the derivative of the RHS, $-e^{-x}$, by the inequality $1 + y \leq e^y$ for all y .




Figure 5.1: A random graph sampled from $\mathcal{G}(n, .01)$ [Von].

two broad categories. First, he designed elaborate randomized constructions of graphs and showed that with nonzero probability, they can possess certain counterintuitive, seemingly impossible properties. This general approach is now called *Ramsey theory*. Second, he showed that for natural random graph models, these graphs – however random – tend to be extremely consistent about certain properties. Today we will study the $\mathcal{G}(n, p)$ random graph, sometimes called Erdős-Rényi graphs based on work by Erdős and Rényi [ER59; ER60]. A random graph from $\mathcal{G}(n, p)$ is an undirected graph over n vertices, where every edge is sampled independently with probability p . By now there is a large catalog of nontrivial and useful properties that, depending on p , are almost certain to appear or not appear in such a graph (for sufficiently large n). Moreover, Erdős and Rényi showed that these properties can vary dramatically with very small changes in p . Consider the following theorem.

Theorem 5.3. *Consider a random graph $G \sim \mathcal{G}(n, p)$ for $p = c/n$, where c is a constant.*

1. *If $c > 1$, then with high probability, there is exactly one connected component of G with $\Omega(n)$ vertices, and all other components have size $\leq O(\log n)$.*
2. *For $c < 1$, then with high probability, all connected components of G has size $< O(\log n)$.*

The parameter c above models the average degree (in expectation). The drama lies in the fact that a tiny change in the average degree c – from .999 to 1.0001 – flips the qualitative nature of a typical random graph from one of many tiny components to essentially one giant component. This is an example of a *threshold phenomena*; alternatively, a *nonlinear dynamic*. Such phenomena is not rare: it occurs in many situations in physics, as well as in models for epidemiology and social networks. Let us briefly mention - without claiming to be very precise - that the sensitivity to c gives some motivation for controlling the “reproductive number” when analyzing and preventing the spread of infectious diseases. The reproductive number is the expected number of healthy individuals that a sick individual effects.

We note that a line of research has obtained a much more refined and detailed understanding than stated in theorem 5.3. We refer the reader to [Bol98, Chapter 7] for further details and other results in this area.

5.2.1 Overview of the proof for $c > 1$

We will prove part 1 of theorem 5.3 in roughly three parts.

Part 1: the gap theorem. Observe that in theorem 5.3 above, regardless of the value of p , there are simply no “medium”-size components, like a component of size $\sqrt{n}$ or of size $n/\log(n)$. The intermediate sizes are ruled out by the following “gap theorem”.

Lemma 5.4. *There is a universal constant $C > 0$, such that for all $\epsilon \in (0, 1)$, and for all $n > 0$ sufficiently large, and $p = (1 + \epsilon)/n$, we have the following. For a random graph $\mathcal{G}(n, p)$, with probability of error $\leq 1/n^2$, no component has k vertices for any value k in the interval*

$$\frac{C \log(n)}{\epsilon^2} \leq k \leq \frac{\epsilon n}{C}.$$

We analyze lemma 5.4 theorem in section 5.3. The proof makes a surprising connection to our discussion on random sums in section 5.1.


Part 2: existence of a large component. Lemma 5.4 establishes that all components are either very small or very big. However it does not assert that there are any big components. The next theorem, proven in section 5.4 and based on analyzing a *Galton-Watson branching process*, shows that any single vertex has a reasonable chance of being in a component that is not small.

Lemma 5.5. *Let $p = (1 + \epsilon)/n$ for $\epsilon > 0$. Let $v \in V$ be a vertex. For all $3 \leq h \leq \epsilon n$, with probability at least $1/h$, v has at least $1 + h$ vertices in its connected component.*

lemma 5.5 implies that there is almost certainly at least one giant component as follows. Let $h = c \log(n)/\epsilon^2$ for a sufficiently large constant c , and let $q = 1/h = \Omega(\epsilon^2/\log(n))$. Call a component “small” if it has at most h vertices. We want to argue that, for $p > 1/n$, there is at least one component that is not small. In conjunction with the gap theorem (lemma 5.4), which rules out all intermediate sizes, this implies that there is at least one giant component of size $\Omega(\epsilon^2 n)$.

By lemma 5.5, any vertex v has at least a probability q of not being in a small component. Now imagine a process where we first randomly select a vertex v and inspect its component. If it is not small, then we have obtained the non-small component we seek. Otherwise, if the component is small, then we throw out v and its component, and randomly select another vertex as v , and repeat. Each vertex we inspect has probability q of not being in a small component. We would have to fail on the order of n/h consecutive samples to conclude there is no small component - which happens with diminishingly small probability. Thus with very high probability, there is at least one component that is not small.

Part 3: uniqueness of the giant component. Can there be two giant components? The answer is no (with high probability) and here is a quick explanation. Instead of sampling from $\mathcal{G}(n, p)$ directly, we can first sample two graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ from $G(n/2, p)$. In the second stage we can sample each cross-edge (v_1, v_2) , where $v_1 \in V_1$ and $v_2 \in V_2$, independently with probability p . Now, by applying the theory we have already developed to G_1 and G_2 , G_1 and G_2 will have some giant components, each of size $\Omega(\epsilon^2 n)$. Note that each graph can only have $O(1/\epsilon^2)$ of them.



Let C_1 be a giant component in G_1 and let C_2 be a giant component in G_2 . We can sample up to $|C_1||C_2| \geq \Omega(\epsilon^4 n^2)$ edges between C_1 and C_2 . Recalling that p is greater than $1/n$, the odds that all $\Omega(\epsilon^4 n^2)$ edges fail to be sampled is vanishingly small. That is, we almost certainly connect C_1 and C_2 . Since there is a limited number of giant components we will almost certainly connect all of them together. Thus, for $p > (1+\epsilon)/n$ for $\epsilon > 0$, we get a *unique* giant component. This establishes theorem 5.3 for $c > 1$.

5.2.2 $c < 1$

The proof for $c < 1$ is simpler and only requires the ideas underlying lemma 5.4. See section 5.3.1.

5.2.3 Directed graphs

One could naturally ask the same questions for directed graphs. Let $D(n, p)$ denote the distribution over *directed graphs* where every directed edge appears independently with probability p . We might similarly ask for the maximum number of vertices reachable from any component, or the size of the maximum strongly component.

It turns out that the analysis of directed graphs can be largely reduced to undirected graphs, as shown by Karp [Kar90] in the following delightfully simple way.

Theorem 5.6. *Let $G \sim G(n, p)$ and $D \sim G(n, p)$, and fix a vertex v . Then the size of the connected component of v in G , and the number of vertices reachable from v in D , are identically distributed.*

Proof. Let us introduce a second distribution of directed random graphs. Let $B(n, p)$ be the distribution of directed graphs where we sample each *undirected* edge $\{u, v\}$ independently with probability p , and for each sampled edge, add both directions (u, v) and (v, u) to the graph. Clearly, for a fixed vertex v , the size of v 's (undirected) component in $G(n, p)$ is distributed identically to the number of vertices reachable from v in $B(n, p)$. We claim that the number of vertices reachable from v in $B(n, p)$ is identically distributed as in $D(n, p)$. At this point let us simply quote Karp [Kar90, Lemma 1] (with minor changes in notation) whose proof is very elegant.

...To see that the last two random variables are identically distributed, note that the probability spaces $B(n, p)$ and $D(n, p)$ differ in only one respect: a digraph G drawn from $B(n, p)$, arc (u, v) is present if and only if arc (v, u) is present, while, in a digraph D drawn from $D(n, p)$, then the event that (v, u) is present is independent of the event that (u, v) is present. Thus no experiment based on checking for the presence or absence of arcs can distinguish between the two probability spaces unless it checks both an arc and its reversal. But any standard sequential algorithm, such as breadth-first search or depth-first search, for building a search tree containing exactly the vertices reachable from vertex 1, checks for the presence of arc (u, v) only if vertex u is in the search tree and v is not; thus it never checks both an arc and its reversal, and accordingly cannot distinguish $B(n, p)$ from $D(n, p)$.

To summarize the excerpt, standard search algorithms for reachability do not distinguish $B(n, p)$ and $D(n, p)$ anyway, so the number of reachable vertices is identically distributed. $\square$

5.3 A gap in component size

In this section we prove lemma 5.4, which asserts that when $p = (1 + \epsilon)/n$ for a constant $\epsilon > 0$, then with high probability, all components are either very small or very large. Our analysis follows an approach due to Karp [Kar90]. His proof is also described in [Bol98]. We will also reuse some of the ideas in the proof to analyze the $p < 1/n$ in section 5.3.1. We first restate lemma 5.4 for the reader's convenience.

Lemma 5.4. *There is a universal constant $C > 0$, such that for all $\epsilon \in (0, 1)$, and for all $n > 0$ sufficiently large, and $p = (1 + \epsilon)/n$, we have the following. For a random graph $\mathcal{G}(n, p)$, with probability of error $\leq 1/n^2$, no component has k vertices for any*

value k in the interval

$$\frac{C \log(n)}{\epsilon^2} \leq k \leq \frac{\epsilon n}{C}.$$

For a vertex $v \in V$, let $C(v) \subset V$ be the (randomized) component of v . To analyze $C(v)$, we imagine revealing $C(v)$ by a search algorithm. We maintain a collection of vertices known to be connected to v ; initially just $\{v\}$. Each iteration i , starting from v , select a vertex v_i that is known to be in $C(v)$, but has not been explored. Then “explore” v_i by inspecting all of the edges incident to v_i , possibly adding to the collection of vertices known to be connected to v (but not yet explored).

We annotate this process as follows. For $i \in \mathbb{N}$, let

- v_i be the vertex that is explored in the i th iteration (or `nil` if all of $C(v)$ has already been explored).


For each $i \in \mathbb{Z}_{\geq 0}$, let

- A_i be the set of vertices known to be in $C(v)$ after i iterations, and let
- B_i be the set of vertices that have been explored.

For the sake of concreteness, one can imagine processing the v_i ’s in BFS order. Recall that BFS marks each vertex when the vertex first encountered, and if the vertex was unmarked, it is added to a queue. The next vertex visited is drawn from the queue. In terms of BFS, then, A_i is the set of vertices marked after i iterations, and B_i is the set of vertices that have left the queue and been fully processed.

Ultimately, B_i, A_i, v_i are built up incrementally as follows.

1. Initially, we have $B_0 = \emptyset$ and $A_0 = \{v\}$.
2. In the first iteration, set $v_1 = v$, set $B_1 = \{v_1\}$, and set $A_1 = A_0 \cup N(v_1)$, where $N(v_1)$ is the (randomized) neighborhood of v_1 .
3. In the i th iteration, if $B_{i-1} \neq A_{i-1}$, then select (any) $v_i \in A_{i-1} \setminus B_{i-1}$. Set $B_i = B_{i-1} \cup \{v_i\}$ and $A_i = A_{i-1} \cup N(v_i)$. Otherwise we terminate with $C(v) = B_{i-1} = A_{i-1}$.



The process terminates when $B_i = A_i$. But since $B_i \subseteq A_i$ and $|B_i| = i$, this is precisely when $|A_i| = i$. As long as $|A_i| \neq i$, A_{i+1} is generated by taking the union of A_i and a random sample of $V - A_i$ where each vertex is included with probability p . Thus we could have generated the sequence of A_i 's instead by the following equivalent process, which omits any mention of B_i or v_i .

1. Initially set $A_0 = \{v\}$.
2. For each $i \in \mathbb{N}$, let S sample each vertex in $V \setminus A_{i-1}$ independently with probability p and set $A_i = A_{i-1} \cup S$.
3. Let i be the first index such that $|A_i| = i$, and return $C(v) = A_i$.

Fix an iteration i . The alternative (but equivalent) process described above exposes a simple distribution for A_i . For any vertex $x \neq v$, we have $x \notin A_i$ iff x failed to be added in each of the first i rounds, which occurs with probability exactly $(1 - p)^i$. Moreover this event is independent across vertices. Thus $|A_i|$ is distributed exactly as the binomial distribution with $n - 1$ coins and probability $1 - (1 - p)^i$; i.e., $|A_i| \sim \mathcal{B}(n - 1, 1 - (1 - p)^i)$.

Lemma 5.7. *Let $i \leq \epsilon n/2((1 + \epsilon))$. Then $\mathbf{E}[|A_i|] \geq (1 + \epsilon)i$.*

Proof. We have

$$(1 - p)^i \leq e^{-ip} \leq 1 - ip + (ip)^2 \leq 1 - ip + \epsilon ip/2 = 1 - (1 - \epsilon/4)ip.$$

where (a) is because $ip = (1 + \epsilon)i/n \leq \epsilon/4$. Thus

$$\mathbf{E}[|A_i|] = 1 + \left(1 - (1 - p)^i\right)(n - 1) \geq (1 - \epsilon/4)ipn \geq (1 + \epsilon)i.$$

□

Lemma 5.8. *Let $i \leq \epsilon n/2(1 + \epsilon)$. Then $\mathbf{P}[|A_i| \leq i] \leq e^{-\epsilon^2 i/8}$.*

Proof. We have

$$\mathbf{P}[|A_i| \leq i] \stackrel{(a)}{\leq} \mathbf{P}[|A_i| \leq (1 - \epsilon/2) \mathbf{E}[|A_i|]] \stackrel{(b)}{\leq} e^{-\epsilon^2 i/8}.$$

Here (a) is by lemma 5.7. (b) is by the tail inequality on binomial distributions, lemma 5.1. □

Let $I = \{i \in \mathbb{N} : 32 \ln(n)/\epsilon^2 \leq i \leq \epsilon n/2((1 + \epsilon))\}$. For all $i \in I$, we have

$$\mathbf{P}[|A_i| \leq i] \leq 1/n^4.$$

By the union bound, we have

$$\mathbf{P}[|A_i| > i \text{ for all } i \in I] \geq 1 - \sum_{i \in I} \mathbf{P}[|A_i| \leq i] \geq 1 - 1/n^3.$$

Thus with probability $\geq 1 - 1/n$, the number of vertices in the connected component of v , $|C(v)|$, does not lie in the range I . Taking the union bound over all $v \in V$ establishes part 1 of lemma 5.4.

5.3.1 Probabilities < 1

Suppose instead that $p = (1 - \epsilon)/n$. Then we have $\mathbf{E}[A_i] \leq (1 - \epsilon/2)i$ unless i is very close to n . In particular, for $i = O(\log(n)/\epsilon^2)$, the probability of $A_i > i$ is $1/\text{poly}(n)$. Thus we see that all components will have size $O(\log(n)/\epsilon^2)$ with high probability, establishing part 2 of lemma 5.4.

5.4 Galton-Watson branching processes for general branch factor

We now move onto the second part of the analysis. By now we have established that there are (with high probability) no “medium” components – all component sizes have either at most $O(\log(n)/\epsilon^2)$ vertices, or at least $\Omega(\epsilon^2 n)$ vertices. Now we want to prove lemma 5.5, which we first restate for the reader’s convenience.

Lemma 5.5. *Let $p = (1 + \epsilon)/n$ for $\epsilon > 0$. Let $v \in V$ be a vertex. For all $3 \leq h \leq \epsilon n$, with probability at least $1/h$, v has at least $1 + h$ vertices in its connected component.*

The proof is by relation to the so-called *Galton-Watson process* that arises in the study of reproducing populations. In the simplest case, imagine a population of size 1. Each generation, each member of the current generation flips k coins, each of which flips heads with probability $1/k$. For each heads, we generate another member of the next generation. The probabilities and number of coins are configured so that each member expects to have one child.

What is the probability that the population survives for h iterations, for a given parameter h ? This is answered by the following.

Theorem 5.9. *Let T be a complete k -ary tree of height h , and suppose every edge is deleted independently with probability at most $1 - 1/k$. Then the probability that there is a leaf connected to the root is $\geq 1/h$ for $h \geq 3$, and $\geq (1 - e^{-1})^h$ for $h \leq 2$.*




Figure 5.2: A complete binary tree of height 3, where each edge was deleted with probability $1/2$.

An example of the case $k = 2$ is drawn in fig. 5.2.

Proof. For $i \in \mathbb{N}$, let p_i be the probability that a particular node at height i is connected to a subleaf. We have $p_0 = 1$. For a node at height $i + 1$, the probability that there is no path to a leaf via a particular child is

$$1 - \frac{1}{k} + \frac{1}{k}(1 - p_i) = 1 - \frac{p_i}{k}.$$

By independence, we have

$$p_{i+1} = 1 - \left(1 - \frac{p_i}{k}\right)^k.$$

Observe that the RHS is increasing in p_i ; thus to lower bound p_{i+1} , we can substitute any lower bound for p_i . We have

$$\begin{aligned} p_0 &= 1, \\ p_1 &= 1 - (1 - 1/k)^k \geq 1 - e^{-1} \geq .63, \\ p_2 &= 1 - (1 - .63/k)^k \geq 1 - e^{-.63} \geq .467, \\ p_3 &\geq 1 - (1 - .467/k)^k \geq 1 - e^{-.467} \geq .373 \geq 1/3. \end{aligned}$$

We claim by induction on i that $p_i \geq 1/i$ for all $i \geq 3$. The base case $i = 3$ was just proven. For the general case,

$$p_{i+1} \stackrel{(a)}{\geq} 1 - (1 - 1/ik)^k \geq 1 - e^{-1/i} \stackrel{(b)}{\geq} \frac{1}{i} - \frac{1}{2i^2} \geq \frac{1}{i+1}$$

Here (a) is by induction. (b) applies the inequality $e^x \leq 1 + x + \frac{1}{2}x^2$ for $x \leq 0$. $\square$

5.4.1 Likelihood of small components

We can use the above branching process to analyze the probability that a given vertex v is in a component of size $\leq h$, for any $h \leq \epsilon n/(1 + \epsilon)$. Recall the sets $B_0, B_1, B_2, \dots$ from section 5.3. Given that $|B_i| \leq h$, we can think of $B_{i+1} - B_i$ as adding (at least) $n/(1 + \epsilon)$ children each with probability $p = (1 + \epsilon)/n$. Either we find new elements for all h rounds - which forces $|B_h| \geq h$ - or we hit $|B_i| = h$ at some point $i < h$. Thus the odds of v acquiring h vertices in its connected component is at least the odds produced by theorem 5.9 for this value of h ; namely, $1/h$. This gives us lemma 5.5.

5.4.2 A more general analysis

One can consider a more general model as follows. Let $X \in \mathbb{Z}_{\geq 0}$ be a random variable taking nonnegative integer values. For each node at a generation i , we sample an independent copy of X , and generate this many children in the next generation. Here we have the following remarkably precise theorem.

Theorem 5.10. *Let n_i denote the number of children in the i th generation in the process described above.*

1. *If $\mathbf{E}[X] < 1$, then $\lim_{i \rightarrow \infty} \mathbf{P}[n_i = 0] = 1$.*
2. *If $\mathbf{E}[X] > 1$, then $\lim_{i \rightarrow \infty} \mathbf{P}[n_i = 0] = q$, where q is the unique solution to $x = \sum_{i \geq 0} \mathbf{P}[X = i]x^i$.*

See the lecture notes by Sinclair [Sin20] for a proof.

5.5 Exercises

Exercise 5.1. Prove lemma 5.2. (Here the important part is not the constant, $1/3$ – any constant $c > 0$ is already interesting.)