Chapter 6
Randomization

Algorithm Theory
WS 2012/13

Fabian Kuhn

Number of Cuts

$\alpha \geq 1$

Theorem: The number of edge cuts of size at most $\alpha \cdot \lambda(G)$ in an n-node graph G is at most $n^{2\alpha}$.

edge conn.

Proof:

using rand. contraction alg $\Longrightarrow$ count cuts

$A \cup B = V$

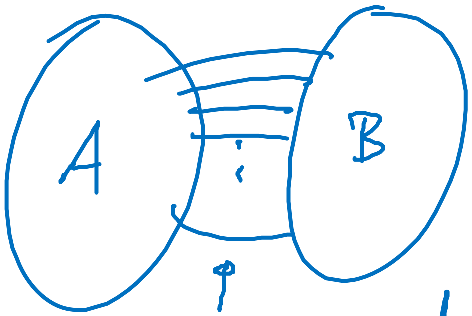

#edges = size of cut (A, B)

Resilience To Edge Failures

- Consider a network (a graph) G with n nodes
- Assume that each link (edge) of G fails independently with probability p
- How large can p be such that the remaining graph is still connected with probability $1 - \varepsilon$?

Chernoff Bounds

$X_i \in \{0,1\}$

$p_i = p$ for all i
$X \sim Bin(n,p)$

- Let $X_1, \dots, X_n$ be independent 0-1 random variables and define $p_i \coloneqq \mathbb{P}(X_i = 1)$. $\Rightarrow \mathbb{E}[X_i] = p_i$
- Consider the random variable $X = \sum_{i=1}^{n} X_i$ # of ones
- We have $\mu \coloneqq \mathbb{E}[X] = \sum_{i=1}^{n} \mathbb{E}[X_i] = \sum_{i=1}^{n} p_i$

Chernoff Bound (Lower Tail):

$$\forall \delta > 0: \ \mathbb{P}(X < (1-\delta)\mu) < e^{-\delta^2 \mu / 2}$$

Chernoff Bound (Upper Tail):

$$\forall \delta > 0: \ \mathbb{P}(X > (1+\delta)\mu) < \left(\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right)^{\mu} < e^{-\delta^2 \mu / 3}$$

holds for $\delta \leq 1$

Chernoff Bounds, Example

$\mathbb{P}(X<(1-\delta)\mu)< e^{-\delta^2\mu/2}$

for $\delta\le 1$: $\mathbb{P}(X>(1+\delta)\mu)< e^{-\delta^2\mu/3}$

Assume that a fair coin is flipped n times. What is the probability to have

1. less than $n/3$ heads?

coin i : rand. X_i, $X_i = \begin{cases}1 & \text{coin } i \text{ is heads}\\ 0 & \text{otherwise}\end{cases}$ $\quad p_i = \mathbb{P}(X_i=1)=1/2$ $\quad$ #heads $X=\sum_{i=1}^{n} X_i$

$\mu = \sum p_i = n\cdot 1/2 = n/2$ $\quad X \sim \mathrm{Bin}(n, 1/2)$

$\mathbb{P}(X<n/3) = \mathbb{P}(X<(1-1/3)n/2) < e^{-(1/3)^2\cdot n/4} = e^{-n/36}$

2. more than $0.51n$ tails?

$\mathbb{P}(X>0.51n) = \mathbb{P}(X>(1+0.02)n/2) < e^{-\frac{0.02^2}{6}\cdot n}$

$0.51\cdot n = (1+\delta)n/2$

$\sqrt{n}$ $\quad \delta = 0.02 \le 1$

3. less than $n/2 - c\ln n$ tails? $\sqrt{n}$

$\mathbb{P}(X<\frac{n}{2}-c\ln n) = \mathbb{P}(X<(1-\frac{2c\ln n}{n})\frac{n}{2}) < e^{-\frac{4c^2 n(\ln n)^2}{n^2}\cdot\frac{n}{2}} = e^{-2c^2(\ln n)^2}$

$\sqrt{n}$

$c\cdot\sqrt{n\cdot\ln n}$

Applied to Edge Cut

$$A \overset{q_k}{\equiv} B$$

- Consider an edge cut (A, B) of size $k = \alpha \cdot \lambda(G)$
- Assume that each edge fails with probability $p \leq 1 - \frac{16 \cdot \ln t}{\lambda(G)}$
- Hence each edge survives with probability $q \geq \frac{16 \cdot \ln t}{\lambda(G)}$
- Probability that at least 1 edge crossing (A, B) survives

rand. var. $X_1, \ldots, X_k$ $X_i = 1$ iff edge i survives $\mathbb{P}(X_i = 1) = q$

$X = \sum X_i$ # of surviving edges

$\mu = \mathbb{E}[X] = k \cdot q$

$\mu = k \cdot q = \alpha \cdot \lambda \cdot \frac{16 \cdot \ln t}{\lambda}$

$= 16 \cdot \alpha \cdot \ln t \geq 2$ assumption $t \geq 2$ suffices

$$\mathbb{P}(X < 1) \underset{\mu \geq 2}{\leq} \mathbb{P}(X < \mu/2) = \mathbb{P}\left(X < \left(1 - \tfrac{1}{2}\right)\mu\right)$$

$$< e^{-\frac{1}{4} \cdot \frac{\mu}{2}} = e^{-2\alpha \ln t} = t^{-2\alpha}$$

upper bound on the prob. that all edges of (A, B) fail

Maintaining Connectivity

- A graph $G = (V, E)$ is connected iff every edge cut (A, B) has size at least 1.

no edge across (A, B) → G not conn

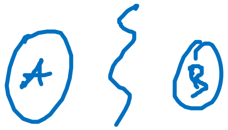

G not connected

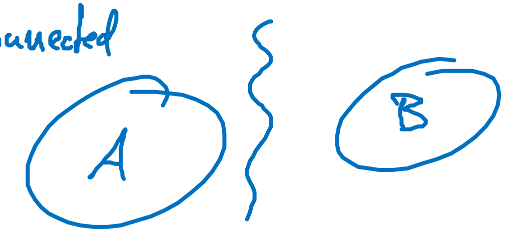

- We need to make sure that every cut keeps at least 1 edge

Maintaining All Cuts of a Certain Size

- The number of cuts of size $k = \alpha\lambda(G)$ is at most $n^{2\alpha}$.

t

Claim: If each edge survives with probability $q \geq \frac{16 \cdot \ln(\beta n)}{\lambda(G)}$, with probability at least $1 - \beta^{-2\alpha}$, at least one edge of each cut of size $k = \alpha\lambda(G)$ survives.

specific cut (A,B) of size k $\Longrightarrow$ $\mathbb{P}(\text{no edge survives}) < t^{-2\alpha} = \beta^{-2\alpha} \cdot n^{-2\alpha}$

$\#$ cuts of size k is x

$x \leq n^{2\alpha}$

$\mathcal{E}_i$: event that cut i does not survive

$i \in \{1, \ldots, x\}$

$\Longrightarrow \mathbb{P}(\mathcal{E}_i) < \beta^{-2\alpha} n^{-2\alpha}$

$$\mathbb{P}(\mathcal{E}_1 \cup \mathcal{E}_2 \cup \ldots \cup \mathcal{E}_x) \leq \sum_{i=1}^{x} \mathbb{P}(\mathcal{E}_i) < x \cdot n^{-2\alpha} \cdot \beta^{-2\alpha} \leq \beta^{-2\alpha}$$

union bound

$\mathbb{P}(A \cup B) \leq \mathbb{P}(A) + \mathbb{P}(B)$

Maintaining All Cuts of a Certain Size

- The number of cuts of size $k = \alpha\lambda(G)$ is at most $n^{2\alpha}$.

Claim: If each edge survives with probability $q \geq \frac{16 \cdot \ln(\beta n)}{\lambda(G)}$, with probability at least $1 - \beta^{-2\alpha}$, at least one edge of each cut of size $k = \alpha\lambda(G)$ survives.

Maintaining Connectivity

$$q = \frac{16 \ln(\beta n)}{\lambda}$$

Theorem: If each edge of a graph G independently fails with probability at most $1 - \underbrace{\frac{8(c+4)\cdot \ln n}{\lambda(G)}}_{q}$, the remaining graph is connected with probability at least $1 - \frac{1}{n^c}$.

Proof: A_k : there is some cut of size $k = \alpha \cdot \lambda(G)$ that does not survive

$$\mathbb{P}(A_k) < \beta^{-2\alpha} \leq \beta^{-2}$$

$$\mathbb{P}(A_\lambda \cup A_{\lambda+1} \cup \dots \cup A_{n^2}) \leq \sum_{k=\lambda}^{n^2} \mathbb{P}(A_k)$$

$$\leq \underline{n^2 \cdot \beta^{-2}} \overset{!}{\leq} n^{-c}$$

$$\beta^2 \geq n^{c+2} \Rightarrow \beta = n^{\frac{c+2}{2}}$$

$$q = \frac{16 \cdot \ln(n^{\frac{c+4}{2}})}{\lambda} = \frac{8(c+4)\ln(n)}{\lambda}$$

If all edges survive with prob. $\Omega\left(\frac{\log(n)}{\lambda(G)}\right)$ then the remaining graph is conn. w.h.p.

Quicksort: High Probability Bound

- To conclude the randomization chapter, let's look at randomized quicksort again

- We have seen that the number of comparisons of randomized quicksort is $O(n \log n)$ in expectation. $2n\ln(n)$

- Can we also show that the number of comparisons is $O(n \log n)$ with high probability? $1 - 1/n$

- **Recall:**

  On each recursion level, each pivot is compared once with each other element that is still in the same "part"

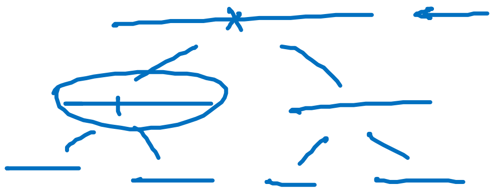

← all element with the pivot

Counting Number of Comparisons

- We looked at 2 ways to count the number of comparisons
  - recursive characterization of the expected number
  - number of different pairs of values that are compared

Let’s consider yet another way:

- Each comparison is between a pivot and a non-pivot
- How many times is a specific array element x compared as a non-pivot?

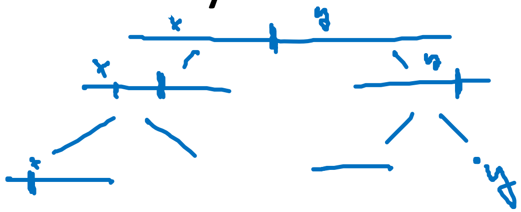

Value x is compared as a non-pivot to a pivot once in every recursion level until one of the following two conditions apply:

1. x is chosen as a pivot
2. x is alone

Successful Recursion Level

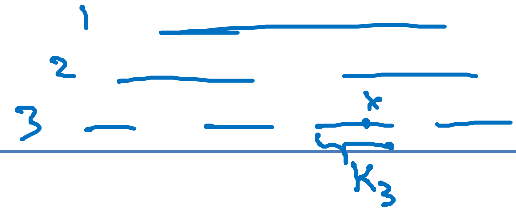

- Consider a specific recursion level ℓ

$K_1 = n$

- Assume that at the beginning of recursion level ℓ, element x is in a sub-array of length K_ℓ that still needs to be sorted.

- If x has been chosen as a pivot before level ℓ, we set $K_\ell := 1$

Definition: We say that recursion level ℓ is successful for element x iff the following is true:

$$K_{\ell+1} = 1 \quad \text{or} \quad K_{\ell+1} \leq \frac{2}{3} \cdot K_\ell$$

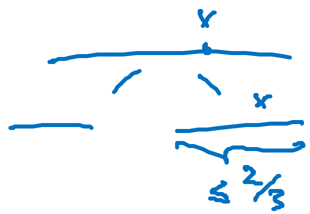

Successful Recursion Level

Lemma: For every recursion level ℓ and every array element x, it holds that level ℓ is successful for x with probability at least $1/3$, independently of what happens in other recursion levels.

Proof:

$K_\ell > 1$

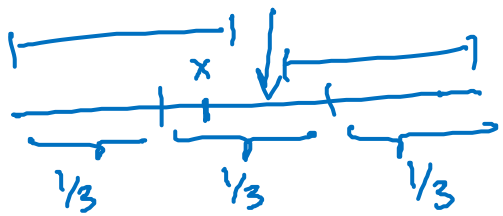

if $K_\ell = 1 \Rightarrow K_{\ell+1} = 1$

$\Rightarrow$ level ℓ is successful

if pivot is in the middle part

$\Rightarrow$ both rem. parts have size $\leq \frac{2}{3} \cdot K_\ell$

$\Rightarrow$ probability for this: $\geq 1/3$

Number of Successful Recursion Levels

Lemma: If among the first ℓ recursion levels, at least $\log_{3/2}(n)$ are successful for element x, we have $K_{\ell+1} = 1$.

Proof: for contradiction, assume $K_{\ell+1} > 1$

$K_1 = n$ $\quad$ $K_{i+1} \leq K_i$ $\quad$ if level i is successful $i \leq \ell$ $\quad$ $K_{i+1} \leq \frac{2}{3} \cdot K_i$

$$K_{\ell+1} \leq n \cdot \left(\frac{2}{3}\right)^{\#\text{successful levels}} \leq n \underbrace{\left(\frac{2}{3}\right)^{\log_{3/2}(n)}}_{1/n} = 1$$

Number of Comparisons for x

Lemma: For every array element x, with high probability, as a non-pivot, x is compared to a pivot at most $O(\log n)$ times.

Proof:

consider ℓ levels

$X_i = \begin{cases} 1 & \text{if level } i \text{ is successful} \\ 0 & \text{otherwise} \end{cases}$

\# of successful levels

$X := \sum_{i=1}^{\ell} X_i$

$Y_i \leq X_i$, $\mathbb{P}(Y_i = 1) = 1/3$, Y_i are independent $\quad Y = \sum Y_i \leq X$

$\mathbb{P}(\text{less than } \log_{3/2}(n) \text{ succ. levels}) \leq \mathbb{P}(Y < \log_{3/2}(n)) < e^{-\Theta(\log n)}$

$\longrightarrow$ choose $\ell = 3c \cdot \log_{3/2}(n) \implies E[Y] = c \cdot \log_{3/2}(n)$

Number of Comparisons for x

Lemma: For every array element x, with high probability, as a non-pivot, x is compared to a pivot at most $O(\log n)$ times.

Proof:

Number of Comparisons

Theorem: With high probability, the total number of comparisons is at most $\boldsymbol{O(n \log n)}$.

Proof:

union bound over all x

$$\frac{1}{n^c} \rightarrow \frac{1}{n^{c-1}}$$