Chapter 7
Randomization

Algorithm Theory
WS 2019/20

Fabian Kuhn

Randomization

Randomized Algorithm:

- An algorithm that uses (or can use) random coin flips in order to make decisions

We will see: randomization can be a powerful tool to

- Make algorithms faster
- Make algorithms simpler
- Make the analysis simpler
  - Sometimes it's also the opposite…
- Allow to solve problems (efficiently) that cannot be solved (efficiently) without randomization
  - True in some computational models (e.g., for distributed algorithms)
  - Not clear in the standard sequential model

Contention Resolution

A simple starter example (from distributed computing)

- Allows to introduce important concepts
- ... and to repeat some basic probability theory

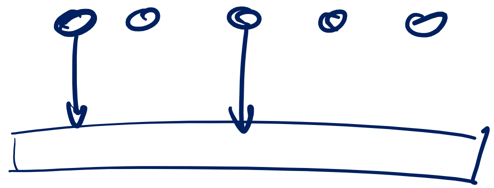

Setting:

- n processes, 1 resource
  (e.g., communication channel, shared database, …)
- There are time slots 1,2,3, …
- In each time slot, only one process can access the resource
- All processes need to regularly access the resource
- If process i tries to access the resource in slot t:
  - Successful iff no other process tries to access the resource in slot t

Algorithm

Algorithm Ideas:

- Accessing the resource deterministically seems hard
  - need to make sure that processes access the resource at different times
  - or at least: often only a single process tries to access the resource
- **Randomized solution:**
  In each time slot, each process tries with probability p.

Analysis:

- How large should p be?
- How long does it take until some process x succeeds?
- How long does it take until all processes succeed?
- What are the probabilistic guarantees?

Analysis

n proc.
time slots $1, \dots, t$

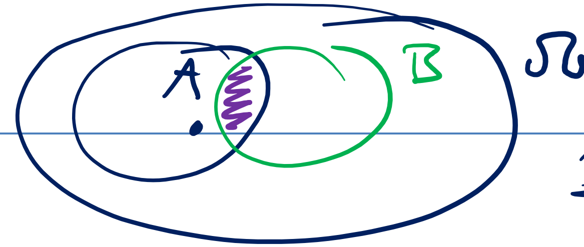

$A \subset \Omega$

Events:

- $\mathcal{A}_{x,t}$: process x tries to access the resource in time slot t
  - Complementary event: $\overline{\mathcal{A}_{x,t}}$

$$\mathbb{P}(\mathcal{A}_{x,t}) = p, \qquad \mathbb{P}(\overline{\mathcal{A}_{x,t}}) = 1 - p$$

- $\mathcal{S}_{x,t}$: process x is successful in time slot t

x successful:
1) x tries to access
2) no other process tries to access

$$\mathcal{S}_{x,t} = \boxed{\mathcal{A}_{x,t}} \cap \left(\bigcap_{y \neq x} \overline{\mathcal{A}_{y,t}} \right)$$

- **Success probability** (for process x):

$$\mathbb{P}(S_{x,t}) = \mathbb{P}(A_{x,t}) \cdot \prod_{y \neq x} \mathbb{P}(\overline{A_{y,t}}) = p \cdot (1-p)^{n-1}$$

choose p s.t. $\mathbb{P}(S_{x,t})$ is maximized

$$p = \frac{1}{n}$$

Fixing p

$\lim_{n\to\infty}\left(1+\frac{1}{n}\right)^n = e$ $\lim_{n\to\infty}\left(1-\frac{1}{n}\right)^n = e^{-1}$

- $\mathbb{P}(\mathcal{S}_{x,t}) = p(1-p)^{n-1}$ is maximized for

$$\boldsymbol{p = \frac{1}{n}} \quad \Longrightarrow \quad \mathbb{P}(\mathcal{S}_{x,t}) = \frac{1}{n}\underbrace{\left(1-\frac{1}{n}\right)^{n-1}}_{\to \frac{1}{e}}.$$

- **Asymptotics:**

$$\text{For } n \geq 2: \quad \frac{1}{4} \leq \left(1-\frac{1}{n}\right)^n < \frac{1}{e} < \left(1-\frac{1}{n}\right)^{n-1} \leq \frac{1}{2}$$

- **Success probability:**

$$\boldsymbol{\frac{1}{en} < \mathbb{P}(\mathcal{S}_{x,t}) \leq \frac{1}{2n}}$$

Time Until First Success $q := \mathbb{P}(S_{x,t}) = \frac{1}{n}\left(1-\frac{1}{n}\right)^{n-1}$

Random Variable T_i:

- $T_i = t$ if proc. i is successful in slot t for the first time

- **Distribution:**

$\mathbb{P}(T_i=1)=q$, $\mathbb{P}(T_i=2)=(1-q)\cdot q$, ..., $\mathbb{P}(T_i=t)=(1-q)^{t-1}\cdot q$

What is $\mathbb{E}[T_i]$?

- T_i is geometrically distributed with parameter

$$q = \mathbb{P}(\mathcal{S}_{i,t}) = \frac{1}{n}\left(1-\frac{1}{n}\right)^{n-1} > \frac{1}{en}.$$

$\frac{1}{en} < q \le \frac{1}{2n}$

- **Expected time** until first success:

$\mathbb{E}[T_i] := \sum_{t=1}^{\infty} t\cdot\mathbb{P}(T_i=t)$

$$\mathbb{E}[T_i] = \frac{1}{q} < en$$

Time Until First Success

Failure Event $\mathcal{F}_{x,t}$: Process x does not succeed in time slots $1, \dots, t$

$$\mathcal{F}_{x,t} = \bigcap_{r=1}^{t} \overline{\mathcal{S}_{x,r}}$$

- The events $\mathcal{S}_{x,t}$ are independent for different t:

$$\mathbb{P}(\mathcal{F}_{x,t}) = \mathbb{P}\left(\bigcap_{r=1}^{t} \overline{\mathcal{S}_{x,r}}\right) = \prod_{r=1}^{t} \mathbb{P}(\overline{\mathcal{S}_{x,r}}) = \big(1 - \underbrace{\mathbb{P}(\mathcal{S}_{x,r})}_{q}\big)^t = (1-q)^t$$

$$\forall \alpha \in \mathbb{R} : 1 + \alpha \le e^{\alpha}$$

- We know that $\overbrace{\mathbb{P}(\mathcal{S}_{x,r})}^{q} > 1/en$:

$$\mathbb{P}(\mathcal{F}_{x,t}) < \Big(\underbrace{1 - \underbrace{\frac{1}{en}}_{\alpha}}_{e^{-1/en}}\Big)^t < e^{-t/en}$$

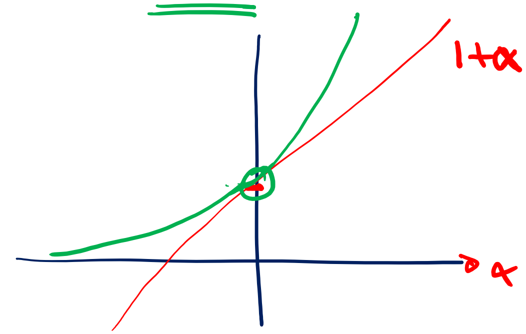

Time Until First Success

$a^{x \cdot y} = (a^x)^y$

No success by time t: $\boxed{\mathbb{P}(\mathcal{F}_{x,t}) < e^{-t/en}}$

$e^{c \ln n} = (e^{\ln n})^c$

$t = \lceil en \rceil$: $\mathbb{P}(\mathcal{F}_{x,t}) < 1/e$

- Generally if $t = \Theta(n)$: constant success probability

$t \geq en \cdot c \cdot \ln n$: $\mathbb{P}(\mathcal{F}_{x,t}) < 1/\underbrace{e^{c \cdot \ln n}}_{e^{-c \cdot \ln n}} = 1/n^c$

- For success probability $1 - 1/n^c$, we need $t = \Theta(n \log n)$.

- We say that i succeeds **with high probability** in $O(n \log n)$ time.

with prob. $\geq 1 - \frac{1}{n^c}$

for any $c > 0$

choice of c only affects the hidden const. in the big-O notation

Time Until All Processes Succeed

Event $\mathcal{F}_t$: some process has not succeeded by time t

$$\mathcal{F}_t = \bigcup_{x=1}^{n} \mathcal{F}_{x,t}$$

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

Union Bound: For events $\mathcal{E}_1, \dots, \mathcal{E}_k$,

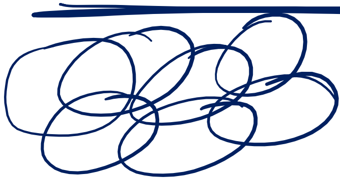

$$\mathbb{P}\left(\bigcup_{x}^{k} \mathcal{E}_x\right) \leq \sum_{x}^{k} \mathbb{P}(\mathcal{E}_x)$$

Probability that not all processes have succeeded by time t:

$$\mathbb{P}(\mathcal{F}_t) = \mathbb{P}\left(\bigcup_{x=1}^{n} \mathcal{F}_{x,t}\right) \underset{\text{union bound}}{\leq} \sum_{x=1}^{n} \mathbb{P}(\mathcal{F}_{x,t}) < n \cdot e^{-t/en}.$$

Time Until All Processes Succeed

Claim: With high probability, all processes succeed in the first $O(n \log n)$ time slots.

Proof:

- $\mathbb{P}(\mathcal{F}_t) < n \cdot e^{-t/en}$
- Set $t = \lceil en \cdot (c+1) \ln n \rceil$

$$\mathbb{P}(\mathcal{F}_t) < n \cdot e^{-\frac{en \cdot (c+1) \ln n}{en}} = n \cdot e^{-(c+1)\ln n} = n \cdot \frac{1}{n^{c+1}} = \frac{1}{n^c}$$

Remark: $\Theta(n \log n)$ time slots are necessary for all processes to succeed with reasonable probability