Probability and Computing – Bounded Differences and Bloom Filters

Stefan Walzer | WS 2025/2026



Content

More Concentration Bounds

Hoeffding: Let $X = X_1 + \cdots + X_n$ where $a_i \leq X_i \leq b_i$, and $X_1, \ldots, X_n$ independent

$\Pr[X - \mathbb{E}[X] \geq t] \leq \exp(-\frac{2t^2}{\sum_{i=1}^{n}(b_i - a_i)^2})$. // proof combines Chernoff with a different lemma

More Concentration Bounds

Hoeffding: Let $X = X_1 + \cdots + X_n$ where $a_i \leq X_i \leq b_i$, and $X_1, \ldots, X_n$ independent

$\Pr[X - \mathbb{E}[X] \geq t] \leq \exp(-\frac{2t^2}{\sum_{i=1}^{n}(b_i - a_i)^2})$. // proof combines Chernoff with a different lemma

Special Case: $X_1, \ldots, X_n$ are Bernoulli random variables

$\Pr[X - \mathbb{E}[X] \geq t] \leq \exp(-2t^2/n)$. // incomparable to the Chernoff bound we saw: $\Pr[X \geq (1+\delta)\mathbb{E}[X]] \leq \left(\frac{e^\delta}{(1+\delta)^{1+\delta}}\right)^\mu$

More Concentration Bounds

Karlsruhe Institute of Technology

Hoeffding: Let $X = X_1 + \cdots + X_n$ where $a_i \leq X_i \leq b_i$, and $X_1, \ldots, X_n$ independent

$\Pr[X - \mathbb{E}[X] \geq t] \leq \exp(-\frac{2t^2}{\sum_{i=1}^n (b_i - a_i)^2})$. // proof combines Chernoff with a different lemma

Special Case: $X_1, \ldots, X_n$ are Bernoulli random variables

$\Pr[X - \mathbb{E}[X] \geq t] \leq \exp(-2t^2/n)$. // incomparable to the Chernoff bound we saw: $\Pr[X \geq (1+\delta)\mathbb{E}[X]] \leq \left(\frac{e^\delta}{(1+\delta)^{1+\delta}}\right)^\mu$

Generalisation: McDiarmid / Method of Bounded Differences

- Assume $X_1, \ldots, X_n$ are independent and $X = f(X_1, \ldots, X_n)$.
- Assume for each i, a change in X_i changes $f(X_1, \ldots, X_n)$ by at most c_i.

$\Pr[X - \mathbb{E}[X] \geq t] \leq \exp(-\frac{2t^2}{\sum_{i=1}^n c_i^2})$ // same bound holds for $\Pr[\mathbb{E}[X] - X \geq t]$.

More Concentration Bounds

Karlsruhe Institute of Technology

Hoeffding: Let $X = X_1 + \cdots + X_n$ where $a_i \le X_i \le b_i$, and $X_1, \ldots, X_n$ independent

$\Pr[X - \mathbb{E}[X] \ge t] \le \exp(-\frac{2t^2}{\sum_{i=1}^n (b_i - a_i)^2})$. // proof combines Chernoff with a different lemma

Special Case: $X_1, \ldots, X_n$ are Bernoulli random variables

$\Pr[X - \mathbb{E}[X] \ge t] \le \exp(-2t^2/n)$. // incomparable to the Chernoff bound we saw: $\Pr[X \ge (1+\delta)\mathbb{E}[X]] \le \left(\frac{e^\delta}{(1+\delta)^{1+\delta}}\right)^\mu$

Generalisation: McDiarmid / Method of Bounded Differences

- Assume $X_1, \ldots, X_n$ are independent and $X = f(X_1, \ldots, X_n)$.
- Assume for each i, a change in X_i changes $f(X_1, \ldots, X_n)$ by at most c_i.

$\Pr[X - \mathbb{E}[X] \ge t] \le \exp(-\frac{2t^2}{\sum_{i=1}^n c_i^2})$ // same bound holds for $\Pr[\mathbb{E}[X] - X \ge t]$.

Proof uses Martingales... not here.

A Simple Use Case for the Method of Bounded Differences

Karlsruhe Institute of Technology

Setting

- fixed graph $G = (V, E)$ and $s, t \in V$
- independent edge lengths $X_1, \ldots, X_m \sim \mathcal{U}([0, 1])$
- let P be the shortest $s - t$ path. // unique with probability 1

McDiarmid / Bounded Differences

- $X_1, \ldots, X_n$ independent and $X = f(X_1, \ldots, X_n)$.
- changing X_i changes $f(X_1, \ldots, X_n)$ by at most c_i.

$\Rightarrow \Pr[X - \mathbb{E}[X] \geq t] \leq \exp(-\frac{2t^2}{\sum_{i=1}^{n} c_i^2})$.

A Simple Use Case for the Method of Bounded Differences

Setting

- fixed graph $G = (V, E)$ and $s, t \in V$
- independent edge lengths $X_1, \ldots, X_m \sim \mathcal{U}([0, 1])$
- let P be the shortest $s - t$ path. // unique with probability 1

McDiarmid / Bounded Differences

- $X_1, \ldots, X_n$ independent and $X = f(X_1, \ldots, X_n)$.
- changing X_i changes $f(X_1, \ldots, X_n)$ by at most c_i.

$\Rightarrow \Pr[X - \mathbb{E}[X] \geq t] \leq \exp(-\frac{2t^2}{\sum_{i=1}^n c_i^2})$.

Example where McDiarmid might be useful

Let $L :=$ length of P.

- $L = f(X_1, \ldots, X_m)$ for independent $X_1, \ldots, X_m$. ✓
- Changing X_i changes L by at most 1. $\rightsquigarrow c_i = 1$. ✓

$\Pr[L - \mathbb{E}[L] \geq t] \leq \exp(-\frac{2t^2}{m})$ // might be useful

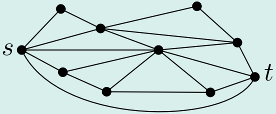

A Simple Use Case for the Method of Bounded Differences

Karlsruhe Institute of Technology

Setting

- fixed graph $G = (V, E)$ and $s, t \in V$
- independent edge lengths $X_1, \ldots, X_m \sim \mathcal{U}([0, 1])$
- let P be the shortest $s - t$ path. // unique with probability 1

McDiarmid / Bounded Differences

- $X_1, \ldots, X_n$ independent and $X = f(X_1, \ldots, X_n)$.
- changing X_i changes $f(X_1, \ldots, X_n)$ by at most c_i.

$\Rightarrow \Pr[X - \mathbb{E}[X] \geq t] \leq \exp(-\frac{2t^2}{\sum_{i=1}^n c_i^2})$.

Example where McDiarmid seems useless

Let $H :=$ number of edges in P. // "hops"

- $H = f(X_1, \ldots, X_m)$ for independent $X_1, \ldots, X_m$. ✓
- Changing X_i might change H by $n - 2$. $\rightsquigarrow c_i = O(n)$.

$\Pr[H - \mathbb{E}[H] \geq t] \leq \exp(-\frac{2t^2}{\mathcal{O}(mn^2)})$ // too weak

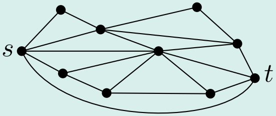

Content

Filter = Approximate Membership Query Data Structure

Karlsruhe Institute of Technology

Setting

- universe D of possible keys
- a set $S \subseteq D$ of $n = |S|$
- a false positive probability ε

Want: Data structure representing S.

Filter = Approximate Membership Query Data Structure

Setting

- universe D of possible keys
- a set $S \subseteq D$ of $n = |S|$
- a false positive probability ε

Want: Data structure representing S.

Space Requirement

- want $\mathcal{O}(n \log_2(1/\varepsilon))$ bits
- *much* smaller than $\mathcal{O}(n \log_2 |D|)$ bits needed for hash table

Filter = Approximate Membership Query Data Structure

Karlsruhe Institute of Technology

Setting

- universe D of possible keys
- a set $S \subseteq D$ of $n = |S|$
- a false positive probability ε

Want: Data structure representing S.

Space Requirement

- want $\mathcal{O}(n \log_2(1/\varepsilon))$ bits
- *much* smaller than $\mathcal{O}(n \log_2 |D|)$ bits needed for hash table

Operations

- **insert** elements to S and **delete** elements from S (optional)
- **query**: given $x \in D$ answer "is $x \in S$?" *approximately*:

$$\textbf{query}(x) = \textbf{YES} \text{ for } x \in S$$

$$\Pr[\textbf{query}(x) = \textbf{NO}] \geq 1 - \varepsilon \text{ for } x \notin S$$

Filter = Approximate Membership Query Data Structure

Karlsruhe Institute of Technology

Setting

- universe D of possible keys
- a set $S \subseteq D$ of $n = |S|$
- a false positive probability ε

Want: Data structure representing S.

Space Requirement

- want $\mathcal{O}(n \log_2(1/\varepsilon))$ bits
- *much* smaller than $\mathcal{O}(n \log_2 |D|)$ bits needed for hash table

Operations

- **insert** elements to S and **delete** elements from S (optional)
- **query**: given $x \in D$ answer "is $x \in S$?" *approximately*:

$$\textbf{query}(x) = \textbf{YES} \text{ for } x \in S$$

$$\Pr[\textbf{query}(x) = \textbf{NO}] \geq 1 - \varepsilon \text{ for } x \notin S$$

a set S:

Anja Blancani
Peter Sanders
Florian Kurpicz
Hape Lehmann
Thomas Worsch

a filter for S:

Anj
Pete
Flo
Hap
Tho

query(Peter Sanders) = **YES**
query(Petra Mutzel) = **YES**
query(Donald Knuth) = **NO**

The **YES** answers are **unreliable**.
The **NO** answers are **reliable**.

Simplified Use Case for Filters

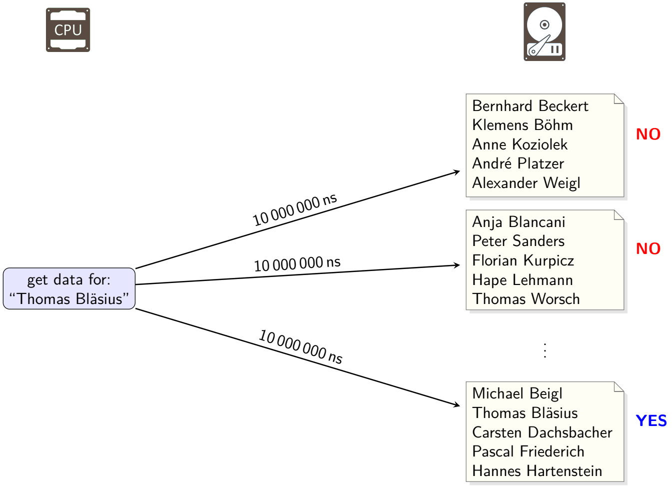

Simplified Use Case for Filters

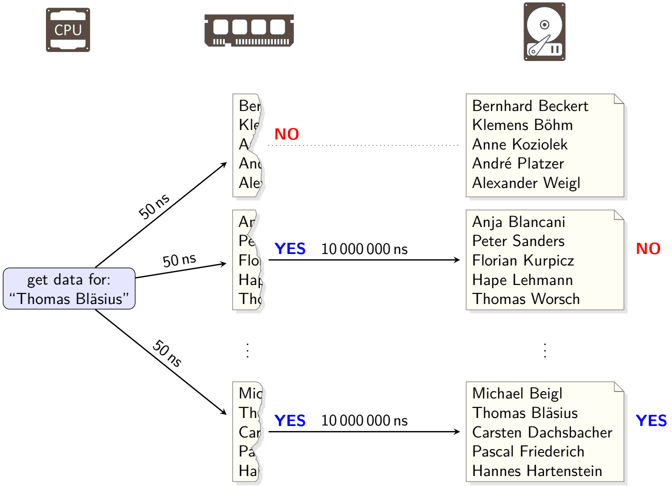

Simplified Use Case for Filters

Karlsruhe Institute of Technology

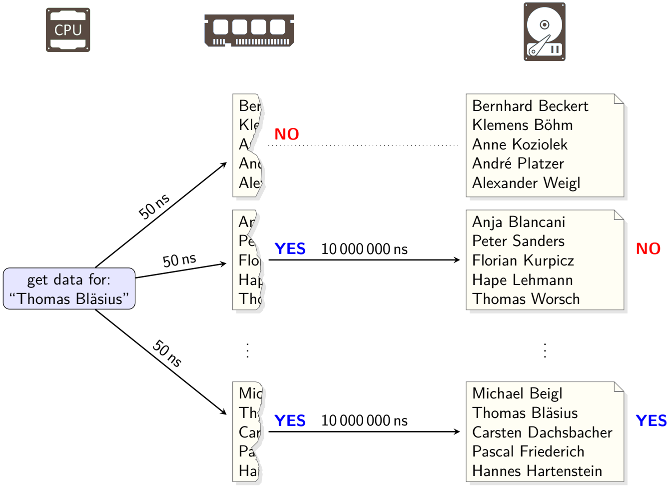

General Idea

If the reliable **NO** answers are frequent, a filter access can replace a (costly) access to a reliable data structure.

Simplified Use Case for Filters

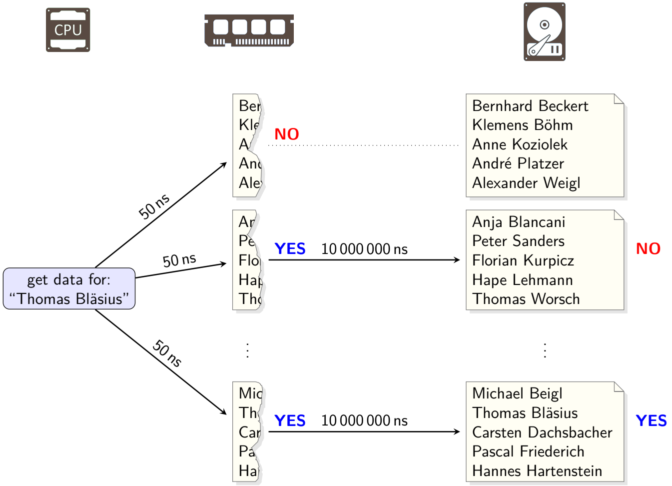

General Idea

If the reliable **NO** answers are frequent, a filter access can replace a (costly) access to a reliable data structure.

Further Example

Filter stores set of malicious URLs.

- Most accessed URLs will be non-malicious.
- Only false positives and true positives have to access reliable data structure (e.g. web-service).

Content

Reminder: SUHA

Simple Uniform Hashing Assumption (SUHA)

- We have access to $h \sim \mathcal{U}(R^D)$ for any R and D.
- h takes $\mathcal{O}(1)$ time to evaluate.
- h takes no space to store.

The Bloom Filter Data Structure

Karlsruhe Institute of Technology

Parameters

m	length of a bit array $A[1..m]$ that we use
$k \in \mathcal{O}(1)$	number of hash functions $h_1, \ldots, h_k \sim \mathcal{U}([m]^D)$
n	number of keys in $S \subseteq D$ (dynamic)
$\alpha \in \mathcal{O}(1)$	load n/m (dynamic)

each cell stores a bit

0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1	2	3	$\cdots$															$\cdots$	m

insert(x):

for $i \in [k]$ **do**
  $A[h_i(x)] = 1$

query(x):

for $i \in [k]$ **do**
  if $A[h_i(x)] = 0$ **then**
    return NO
return YES

delete(x):

⟨ not supported ⟩

The Bloom Filter Data Structure

Parameters

m	length of a bit array $A[1..m]$ that we use
$k \in \mathcal{O}(1)$	number of hash functions $h_1, \ldots, h_k \sim \mathcal{U}([m]^D)$
n	number of keys in $S \subseteq D$ (dynamic)
$\alpha \in \mathcal{O}(1)$	load n/m (dynamic)

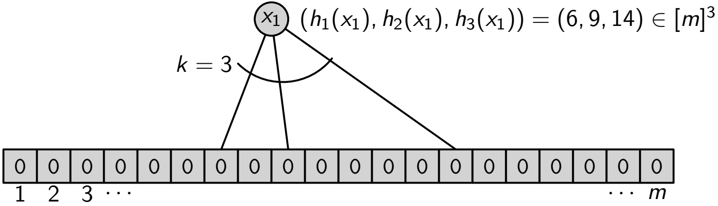

insert(x):

for $i \in [k]$ **do**
  $A[h_i(x)] = 1$

query(x):

for $i \in [k]$ **do**
  if $A[h_i(x)] = 0$ **then**
    return NO
return YES

delete(x):

$\langle$ not supported $\rangle$

The Bloom Filter Data Structure

Karlsruhe Institute of Technology

Parameters

m	length of a bit array $A[1..m]$ that we use
$k \in \mathcal{O}(1)$	number of hash functions $h_1, \ldots, h_k \sim \mathcal{U}([m]^D)$
n	number of keys in $S \subseteq D$ (dynamic)
$\alpha \in \mathcal{O}(1)$	load n/m (dynamic)

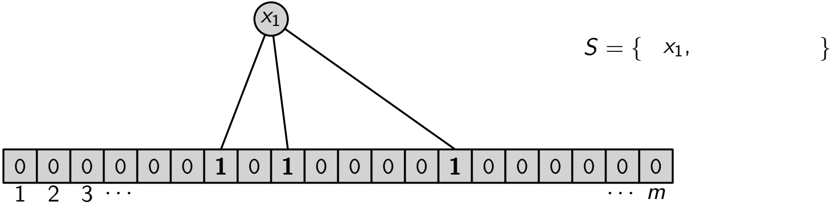

insert(x):

for $i \in [k]$ **do**
  $A[h_i(x)] = 1$

query(x):

for $i \in [k]$ **do**
  if $A[h_i(x)] = 0$ **then**
    return NO

return YES

delete(x):

$\langle$ not supported $\rangle$

The Bloom Filter Data Structure

Parameters

m	length of a bit array $A[1..m]$ that we use
$k \in \mathcal{O}(1)$	number of hash functions $h_1, \ldots, h_k \sim \mathcal{U}([m]^D)$
n	number of keys in $S \subseteq D$ (dynamic)
$\alpha \in \mathcal{O}(1)$	load n/m (dynamic)

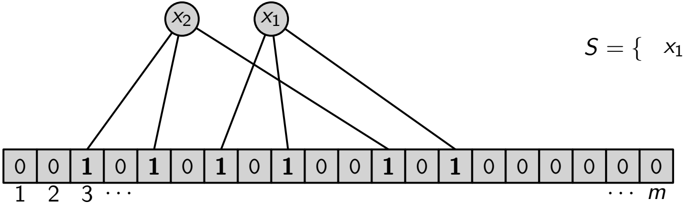

$S = \{\ x_1, x_2, \qquad \}$

insert(x):

for $i \in [k]$ **do**
  $A[h_i(x)] = 1$

query(x):

for $i \in [k]$ **do**
  if $A[h_i(x)] = 0$ **then**
    return NO

return YES

delete(x):

⟨ not supported ⟩

The Bloom Filter Data Structure

Karlsruhe Institute of Technology

Parameters

m	length of a bit array $A[1..m]$ that we use
$k \in \mathcal{O}(1)$	number of hash functions $h_1, \ldots, h_k \sim \mathcal{U}([m]^D)$
n	number of keys in $S \subseteq D$ (dynamic)
$\alpha \in \mathcal{O}(1)$	load n/m (dynamic)

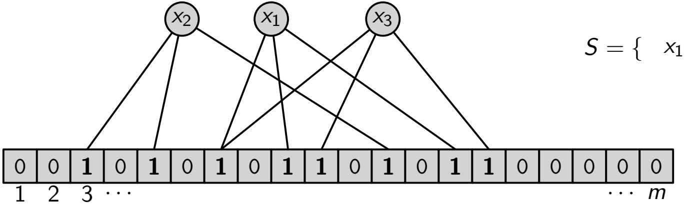

$S = \{ \; x_1, x_2, x_3, \qquad \}$

insert(x):

for $i \in [k]$ **do**
  $A[h_i(x)] = 1$

query(x):

for $i \in [k]$ **do**
  if $A[h_i(x)] = 0$ **then**
    return NO

return YES

delete(x):

$\langle$ not supported $\rangle$

The Bloom Filter Data Structure

Parameters

m	length of a bit array $A[1..m]$ that we use
$k \in \mathcal{O}(1)$	number of hash functions $h_1, \ldots, h_k \sim \mathcal{U}([m]^D)$
n	number of keys in $S \subseteq D$ (dynamic)
$\alpha \in \mathcal{O}(1)$	load n/m (dynamic)

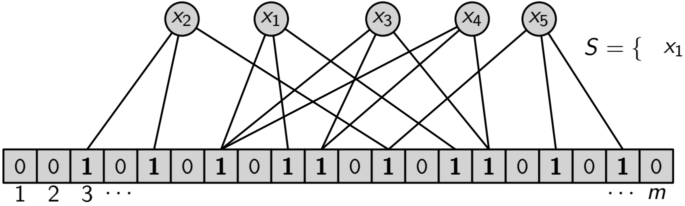

$S = \{ \ x_1, x_2, x_3, x_4, x_5 \}$

insert(x):

for $i \in [k]$ **do**
  $A[h_i(x)] = 1$

query(x):

for $i \in [k]$ **do**
  if $A[h_i(x)] = 0$ **then**
    return NO
return YES

delete(x):

$\langle$ not supported $\rangle$

The Bloom Filter Data Structure

Karlsruhe Institute of Technology

Parameters

m	length of a bit array $A[1..m]$ that we use
$k \in \mathcal{O}(1)$	number of hash functions $h_1, \dots, h_k \sim \mathcal{U}([m]^D)$
n	number of keys in $S \subseteq D$ (dynamic)
$\alpha \in \mathcal{O}(1)$	load n/m (dynamic)

0	0	**1**	0	**1**	0	**1**	0	**1**	**1**	0	**1**	0	**1**	**1**	0	**1**	0	**1**	0
1	2	3	$\cdots$															$\cdots$	m

insert(x):

for $i \in [k]$ **do**
  $A[h_i(x)] = 1$

query(x):

for $i \in [k]$ **do**
  if $A[h_i(x)] = 0$ **then**
    return NO

return YES

delete(x):

$\langle$ not supported $\rangle$

The Bloom Filter Data Structure

Parameters

m	length of a bit array $A[1..m]$ that we use
$k \in \mathcal{O}(1)$	number of hash functions $h_1, \ldots, h_k \sim \mathcal{U}([m]^D)$
n	number of keys in $S \subseteq D$ (dynamic)
$\alpha \in \mathcal{O}(1)$	load n/m (dynamic)

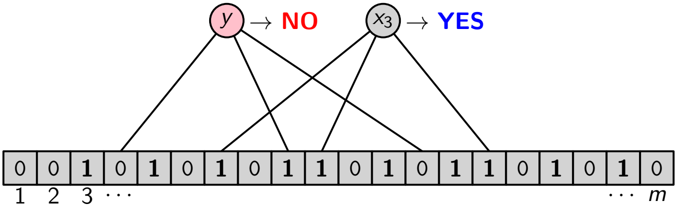

insert(x):

for $i \in [k]$ **do**
  $A[h_i(x)] = 1$

query(x):

for $i \in [k]$ **do**
  if $A[h_i(x)] = 0$ **then**
    return NO
return YES

delete(x):

⟨ not supported ⟩

Content

Preparation

Karlsruhe Institute of Technology

Exercise: Some approximations of e

$$\forall n \in \mathbb{N} : (1+\tfrac{1}{n})^n \leq e \leq (1+\tfrac{1}{n})^{n+1}$$
$$\text{and} \quad (1-\tfrac{1}{n})^n \leq e^{-1} \leq (1-\tfrac{1}{n})^{n-1}.$$

Corollaries

$$\forall n \in \mathbb{N} : (1+\tfrac{1}{n})^n = e - \mathcal{O}(1/n)$$
$$\text{and} \quad (1-\tfrac{1}{n})^n = e^{-1} - \mathcal{O}(1/n).$$

Bloom Filter Analysis (i)

Lemma

Assume $S = \{x_1, \ldots, x_n\}$ is inserted into the Bloom filter. Let $(A_1, \ldots, A_m) \in \{0,1\}^m$ be the random filter state and $Z := \sum_{i=1}^{m}(1 - A_i)$ the number of zeroes. Then

- i $\mathbb{E}[\frac{Z}{m}] = (1 - \frac{1}{m})^{m\alpha k} = e^{-\alpha k} - o(1)$
- ii For $y \notin S$: $\Pr[\textbf{query}(y) = \textbf{YES} \mid Z = z] = (1 - \frac{z}{m})^k$

0	0	**1**	0	**1**	0	**1**	0	**1**	**1**	0	**1**	0	**1**	**1**	0	**1**	0	**1**	0
1	2	3	$\cdots$															$\cdots$	m

Bloom Filter Analysis (i)

Lemma

Assume $S = \{x_1, \dots, x_n\}$ is inserted into the Bloom filter. Let $(A_1, \dots, A_m) \in \{0,1\}^m$ be the random filter state and $Z := \sum_{i=1}^m (1 - A_i)$ the number of zeroes. Then

- i $\mathbb{E}[\frac{Z}{m}] = (1 - \frac{1}{m})^{m\alpha k} = e^{-\alpha k} - o(1)$
- ii For $y \notin S$: $\Pr[\textbf{query}(y) = \textbf{YES} \mid Z = z] = (1 - \frac{z}{m})^k$

0	0	**1**	0	**1**	0	**1**	0	**1**	**1**	0	**1**	0	**1**	**1**	0	**1**	0	**1**	0
1	2	3 ···																···	m

Proof of (i).

$$\mathbb{E}[\tfrac{Z}{m}] = \tfrac{1}{m}\mathbb{E}[\sum_{i=1}^{m}(1 - A_i)] = \frac{1}{m}\sum_{i=1}^{m}\Pr[A_i = 0] = \frac{1}{m}\sum_{i=1}^{m}\Pr[A_1 = 0] = \Pr[A_1 = 0]$$

$$= \Pr[\forall x \in S : \forall i \in [k] : h_i(x) \neq 1] \overset{\text{SUHA}}{=} \prod_{x \in S}\prod_{i \in [k]}\Pr[h_i(x) \neq 1] \overset{\text{SUHA}}{=} \prod_{x \in S}\prod_{i \in [k]}(1 - \tfrac{1}{m})$$

$$= (1 - \tfrac{1}{m})^{nk} = (1 - \tfrac{1}{m})^{m\alpha k} = (e^{-1} - o(1))^{\alpha k} = e^{-\alpha k} - o(1).$$

Bloom Filter Analysis (i)

Karlsruhe Institute of Technology

Lemma

Assume $S = \{x_1, \dots, x_n\}$ is inserted into the Bloom filter. Let $(A_1, \dots, A_m) \in \{0,1\}^m$ be the random filter state and $Z := \sum_{i=1}^{m}(1 - A_i)$ the number of zeroes. Then

- i $\mathbb{E}[\frac{Z}{m}] = (1 - \frac{1}{m})^{m\alpha k} = e^{-\alpha k} - o(1)$
- ii For $y \notin S$: $\Pr[\textbf{query}(y) = \textbf{YES} \mid Z = z] = (1 - \frac{z}{m})^k$

0	0	**1**	0	**1**	0	**1**	0	**1**	**1**	0	**1**	0	**1**	**1**	0	**1**	0	**1**	0
1	2	3	⋯															⋯	m

Proof of (ii).

$$\Pr[\textbf{query}(y) = \textbf{YES} \mid Z = z] = \Pr[\forall i \in [k] : A_{h_i(y)} = 1 \mid Z = z] \overset{\text{SUHA}}{=} \prod_{i \in [k]} \left(\frac{m-z}{m}\right) = (1 - \tfrac{z}{m})^k.$$

How should a Bloom filter be configured?

Karlsruhe Institute of Technology

Approximate false positive rate

From the previous Lemma we get for $y \notin S$:

$$\varepsilon = \Pr[\mathbf{query}(y) = \mathbf{YES}] \approx \Pr[\mathbf{query}(y) = \mathbf{YES} \mid Z = \mathbb{E}[Z]]$$
$$\overset{\text{ii}}{=} \left(1 - \frac{\mathbb{E}[Z]}{m}\right)^k \overset{\text{i}}{=} (1 - e^{-\alpha k} + o(1))^k \approx (1 - e^{-\alpha k})^k.$$

How should a Bloom filter be configured?

Approximate false positive rate

From the previous Lemma we get for $y \notin S$:

$$\varepsilon = \Pr[\textbf{query}(y) = \textbf{YES}] \approx \Pr[\textbf{query}(y) = \textbf{YES} \mid Z = \mathbb{E}[Z]]$$

$$\overset{\text{ii}}{=} \left(1 - \frac{\mathbb{E}[Z]}{m}\right)^k \overset{\text{i}}{=} (1 - e^{-\alpha k} + o(1))^k \approx (1 - e^{-\alpha k})^k.$$

Which k minimises ε? (when α is fixed)

$$\begin{aligned}
&\arg\min_{k \in \mathbb{N}} (1 - e^{-\alpha k})^k \\
= &\arg\min_{k \in \mathbb{N}} (1 - e^{-\alpha k})^{\alpha k} \\
\approx &\frac{1}{\alpha} \arg\min_{t \in \mathbb{R}_+} (1 - e^{-t})^t \\
= &\frac{1}{\alpha} \arg\min_{t \in \mathbb{R}_+} t \ln(1 - e^{-t})
\end{aligned}$$

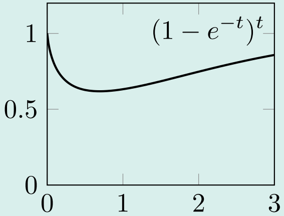

- plot $(1 - e^{-t})^t \rightsquigarrow$ one global minimum.
- deriving $t \ln(1 - e^{-t})$ gives $\ln(1 - e^{-t}) + \frac{te^{-t}}{1-e^{-t}}$
- $t = \ln(2)$ is root of the derivative.

$\hookrightarrow$ $k = \ln(2)/\alpha$ is optimal for fixed α.
$\hookrightarrow$ choose α and k such that $\alpha k = \ln(2)$

False Positive Probability and Space

Karlsruhe Institute of Technology

Intuition for optimality of $\alpha k = \ln(2)$

- gives $\mathbb{E}[\frac{Z}{m}] \approx e^{-\alpha k} = \frac{1}{2}$
- maximises *entropy* of the filter bits

Theorem

A Bloom filter with $k \in \mathbb{N}$ hash functions and load factor $\alpha = \ln(2)/k$ has

- i a **space requirement** of $m = n/\alpha = \frac{kn}{\ln 2} \approx 1.44kn$ bits
- ii and a **false positive probability** of $\varepsilon = 2^{-k} + o(1)$.

(i) Space Requirement

Nothing to proof.

(ii) False Positive Probability

- **Know:** $\mathbb{E}[\frac{Z}{m}] \approx e^{-\alpha k} = \frac{1}{2}$ by choice of α
- **Know:** $\frac{Z}{m} = \frac{1}{2} \Rightarrow \varepsilon = 2^{-k}$
- **Need:** $\frac{Z}{m} \approx \mathbb{E}[\frac{Z}{m}]$, i.e. good concentration

Concentration bound for Z

Lemma

- i $\Pr[Z \leq \mathbb{E}[Z] - t \quad] \leq \exp(-\Theta(t^2/m))$ for any $t > 0$,
- ii $\Pr[Z \leq \mathbb{E}[Z] - m^{2/3}] \leq \exp(-\Theta(m^{1/3}))$ by setting $t = m^{2/3}$.

Concentration bound for Z

Lemma

- i $\Pr[Z \le \mathbb{E}[Z] - t \quad] \le \exp(-\Theta(t^2/m))$ for any $t > 0$,
- ii $\Pr[Z \le \mathbb{E}[Z] - m^{2/3}] \le \exp(-\Theta(m^{1/3}))$ by setting $t = m^{2/3}$.

Difficulty: Filter bits $A_1, \ldots, A_m$ are not independent

- note e.g. that $A_1 = \ldots = A_{m-1} = 0 \Rightarrow A_m = 1$
- Generally: $A_1, \ldots, A_m$ are *negatively associated*:
  - Knowing that some bits are zero makes it more likely that others are 1 and vice versa.
- Standard Chernoff bounds don't apply.

Concentration bound for Z

Karlsruhe Institute of Technology

Lemma

- i $\Pr[Z \le \mathbb{E}[Z] - t \quad] \le \exp(-\Theta(t^2/m))$ for any $t > 0$,
- ii $\Pr[Z \le \mathbb{E}[Z] - m^{2/3}] \le \exp(-\Theta(m^{1/3}))$ by setting $t = m^{2/3}$.

McDiarmid / Bounded Differences

- $X_1, \ldots, X_n$ independent and $X = f(X_1, \ldots, X_n)$.
- changing X_i changes $f(X_1, \ldots, X_n)$ by at most c_i.

$\Rightarrow \Pr[\mathbb{E}[X] - X \ge t] \le \exp(-\frac{2t^2}{\sum_{i=1}^n c_i^2})$.

Difficulty: Filter bits $A_1, \ldots, A_m$ are not independent

- note e.g. that $A_1 = \ldots = A_{m-1} = 0 \Rightarrow A_m = 1$
- Generally: $A_1, \ldots, A_m$ are *negatively associated*:
  - Knowing that some bits are zero makes it more likely that others are 1 and vice versa.
- Standard Chernoff bounds don't apply.

Concentration bound for Z

Karlsruhe Institute of Technology

Lemma

- i $\Pr[Z \leq \mathbb{E}[Z] - t \quad] \leq \exp(-\Theta(t^2/m))$ for any $t > 0$,
- ii $\Pr[Z \leq \mathbb{E}[Z] - m^{2/3}] \leq \exp(-\Theta(m^{1/3}))$ by setting $t = m^{2/3}$.

McDiarmid / Bounded Differences

- $X_1, \ldots, X_n$ independent and $X = f(X_1, \ldots, X_n)$.
- changing X_i changes $f(X_1, \ldots, X_n)$ by at most c_i.

$\Rightarrow \Pr[\mathbb{E}[X] - X \geq t] \leq \exp(-\frac{2t^2}{\sum_{i=1}^n c_i^2})$.

Proof of (i) using the method of bounded differences.

- Z is a function of kn independent hash values // SUHA
- each hash value can change Z by at most 1

$$\Rightarrow \Pr[Z \leq \mathbb{E}[Z] - t] \leq \Pr[\mathbb{E}[Z] - Z \geq t] = \exp\left(\frac{-2t^2}{nk}\right) = \exp\left(\frac{-2t^2}{m\alpha k}\right) = \exp\left(\frac{-2t^2}{m\ln(2)}\right). \quad \square$$

False Positive Probability of Bloom filters

Karlsruhe Institute of Technology

Proof of Theorem, part (ii)

By choice of k and α we have $\mathbb{E}[\frac{Z}{m}] = e^{-\alpha k} - o(1) = \frac{1}{2} - o(1)$.
Let $y \notin S$ and $B = \lfloor \mathbb{E}[Z] - m^{2/3} \rfloor$.

$$\varepsilon = \Pr[\textbf{query}(y) = \textbf{YES}] \overset{\text{LTP}}{=} \sum_{z=1}^{m} \Pr[Z = z] \cdot \Pr[\textbf{query}(y) = \textbf{YES} \mid Z = z] = \sum_{z=1}^{m} \Pr[Z = z] \cdot (1 - \tfrac{z}{m})^k$$

$$\leq \sum_{z=1}^{B} \Pr[Z = z] + \sum_{z=B+1}^{m} \Pr[Z = z](1 - \tfrac{B+1}{m})^k \leq \Pr[Z \leq B] + (1 - \tfrac{B+1}{m})^k$$

$$\leq \Pr[Z \leq \mathbb{E}[Z] - m^{2/3}] + \left(1 - \tfrac{\mathbb{E}[Z]-m^{2/3}}{m}\right)^k \overset{\text{ii}}{\leq} \exp(-\Theta(m^{1/3})) + (1 - \tfrac{1}{2} + o(1))^k = 2^{-k} + o(1). \quad \square$$

How to Configure Your Bloom Filter

Karlsruhe Institute of Technology

Theorem

A Bloom filter with $k \in \mathbb{N}$ hash functions and load factor $\alpha = \ln(2)/k$ has

i. a **space requirement** of $m = n/\alpha = \frac{kn}{\ln 2} \approx 1.44kn$ bits
ii. and a **false positive probability** of $\varepsilon = 2^{-k} + o(1)$.

How to determine *m* and *k* (the parameters you actually need)

1. n: determined by input
2. ε: choose a trade-off between space usage and false positive probability
   - If utility comes from negative answers "$x \notin S$, definitely" and running time is negligible, then:
     - want to maximise utility − disutility, where: ($\propto$ means "proportional to")
     - utility $\propto$ negative answers $=$ queries $\cdot \Pr_{x \sim \text{QueryDistribution}}[x \notin S] \cdot (1 - \varepsilon)$
     - disutility $\propto$ space consumption $= 1.44 \log_2(1/\varepsilon) n$ bits of RAM or cache
3. compute $k = \lceil \log_2(1/\varepsilon) \rceil$ // effectively restricts ε to powers of 2
4. compute $\alpha = \ln(2)/k$ and $m = \lceil n/\alpha \rceil$

Remarks

Karlsruhe Institute of Technology

Much, much more is known

- more functionality
  - $\hookrightarrow$ counting Bloom filters support deletions
- better query times
  - $\hookrightarrow$ blocked Bloom filters improve cache efficiency
- better space efficiency
  - $\hookrightarrow$ cuckoo filters use $n \log_2(1/\varepsilon) + \mathcal{O}(n)$ bits rather than $\approx 1.44 n \log_2(1/\varepsilon)$ bits
  - $\hookrightarrow$ static filters (no insertions or deletions) use $n \log_2(1/\varepsilon) + o(n)$ bits.
- ...

Conclusion

Karlsruhe Institute of Technology

Method of Bounded Differences

- requires $X = f(X_1, \dots, X_n)$ for independent $X_1, \dots, X_n$
- requires that changing X_i changes X only a little
- gives Chernoff-like concentration bounds

Approximate Membership Data Structure

- decides “is $x \in S$?” with *false positive probability* ε
- much smaller than exact membership data structure
  - $\approx 1.44 \log_2(1/\varepsilon)$ bits per element for Bloom filters
  - often fits into cache or RAM
- used to avoid costly access to exact data structure
  - reliable on **NO** answers
  - useful if **NO** answers are frequent

Anhang: Mögliche Prüfungsfragen I

- Methode der beschränkten Differenzen
  - Was sind die Voraussetzungen und Garantien von McDiarmids Ungleichung?
  - Was wäre ein Beispiel wo sie gute bzw. schlechte Garantien liefert?
- Approximate-Membership-Query Datenstrukturen im Allgemeinen
  - Welche Aufgabe hat eine AMQ Datenstruktur?
  - Was ist der Vorteil gegenüber einer exakten Datenstruktur?
  - Was wäre ein Anwendungsfall, in dem eine AMQ Datenstruktur nützlich ist?
- Bloomfilter
  - Wie ist ein Bloomfilter aufgebaut und welche Operationen unterstützt er?
  - Welche Parameter gibt es, und wie hängen diese zusammen?
  - Was hat unsere Analyse zur geschickten Wahl der Parameter zu sagen? Wie werden die übrigen Parameter gewählt? Welcher Speicherverbrauch ergibt sich?
  - Fragen zur Analyse
    - Welche Anzahl von Nullen bzw. Einsen erwarten wir?
    - Wie hängt die falsch-positiv Wahrscheinlichkeit mit der Anzahl Nullen bzw. Einsen zusammen?
    - Wir kann man argumentieren, dass die Anzahl Nullen bzw. Einsen im Bloomfilter nahe am Erwartungswert liegt?