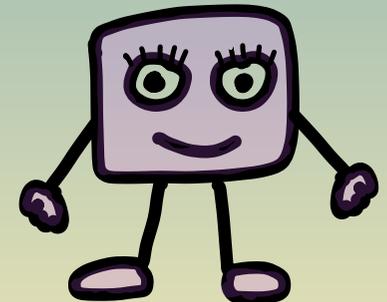


Filters



Michael A. Bender

This is the
advertised
talk title.



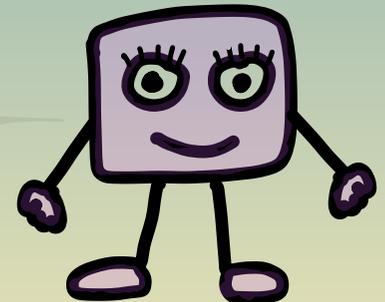
Time To Change Your Filter



*This title
is more
to the point.
(more explanation
later).*

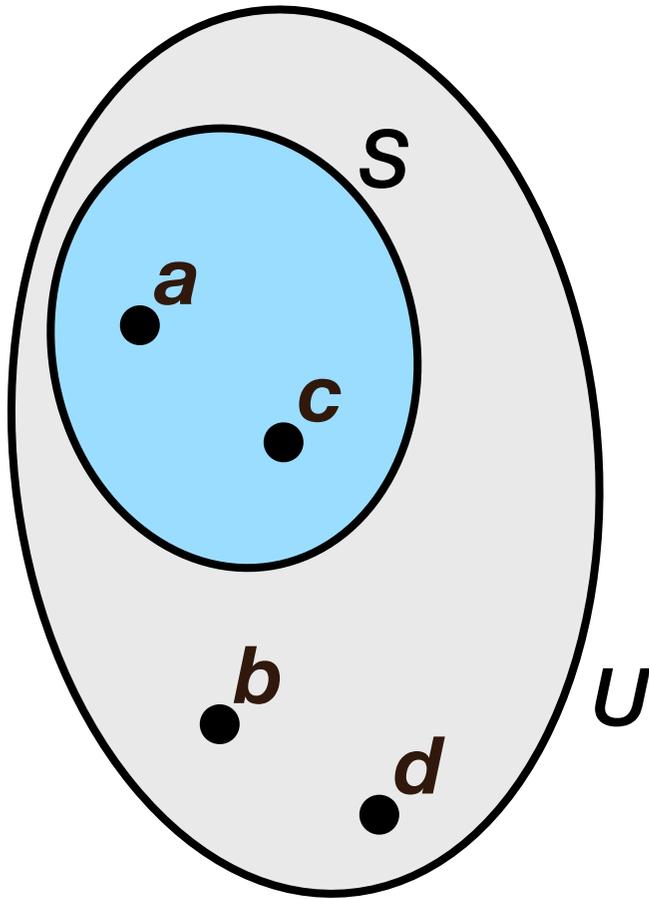
Michael A. Bender

*Now... what's
a filter.*



Dictionary Data Structure

A dictionary maintains a set S from universe U .



member(a): ✓

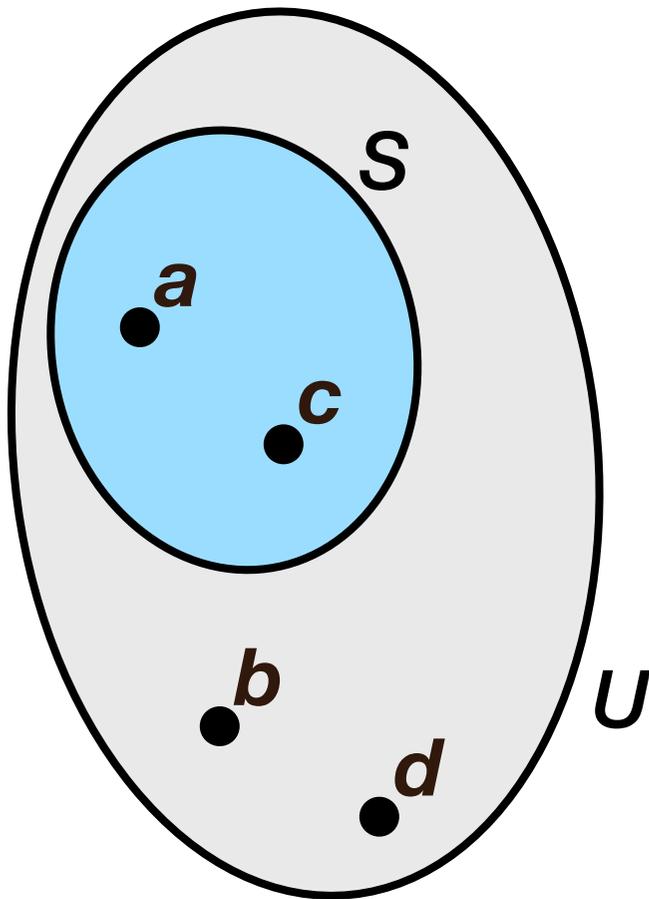
member(b): ✗

member(c): ✓

member(d): ✗

A dictionary supports membership queries on S .

A filter is an *approximate* dictionary.



member(a): ✓

member(b): ✗

member(c): ✓

member(d): ✓



false
positive

A filter supports approximate membership queries on S .

A Filter Guarantees a False-Positive Rate ϵ

if $q \in S$, return  with probability 1 true positive

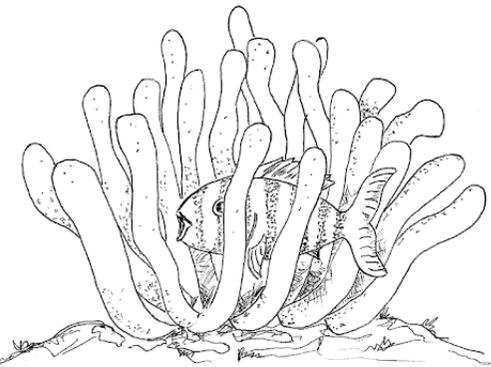
if $q \notin S$, return $\left\{ \begin{array}{l} \times \\ \checkmark \end{array} \right.$ with probability $> 1 - \epsilon$ true negative
with probability $\leq \epsilon$  false positive

↑
one-sided errors

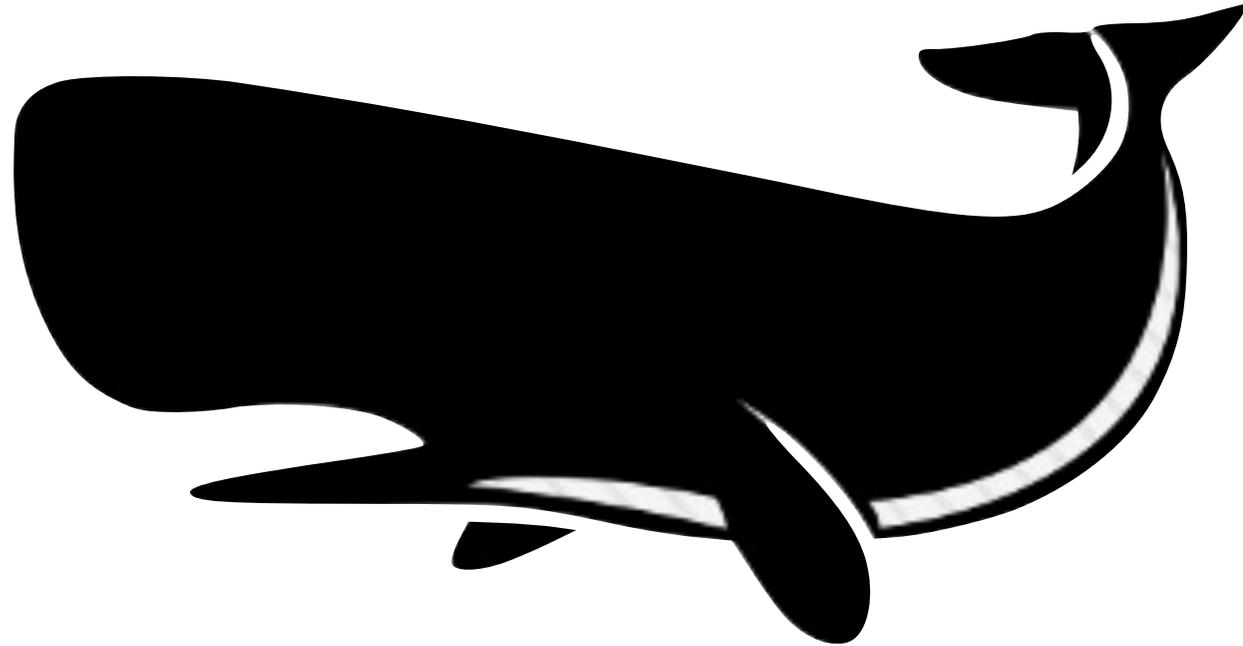
False-positive rate enables filters to be compact

$$\text{space} \geq n \log(\overbrace{1/\varepsilon}^{\text{Small}})$$

$$\text{space} = \Omega(n \log \overbrace{|U|}^{\text{large}})$$



Filter

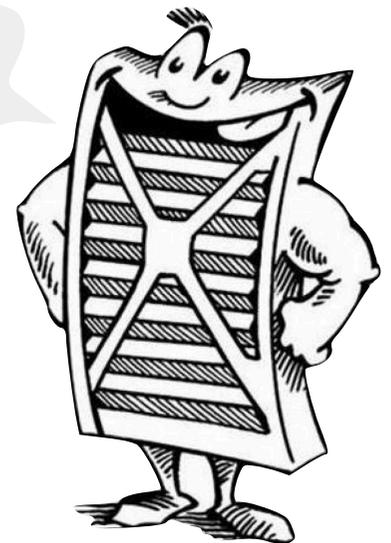


Dictionary

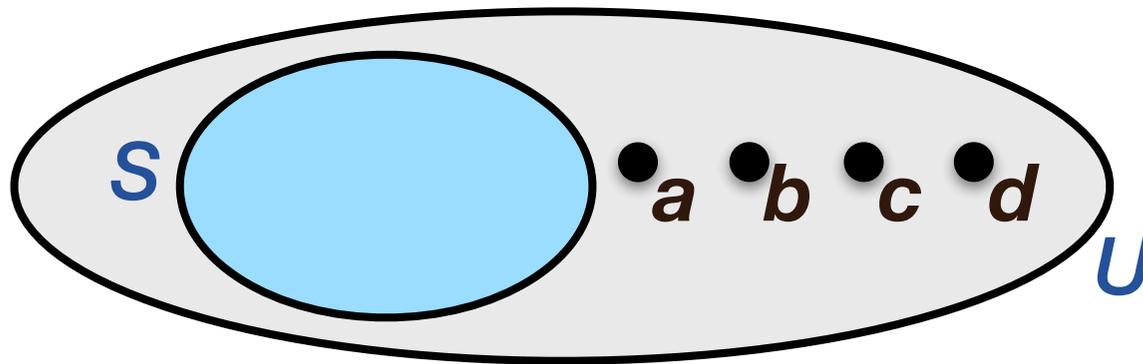
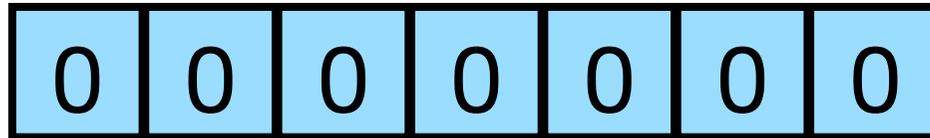
Talk So Far

- Filter data structure
- Next: the **Bloom filter** [Bloom '70]

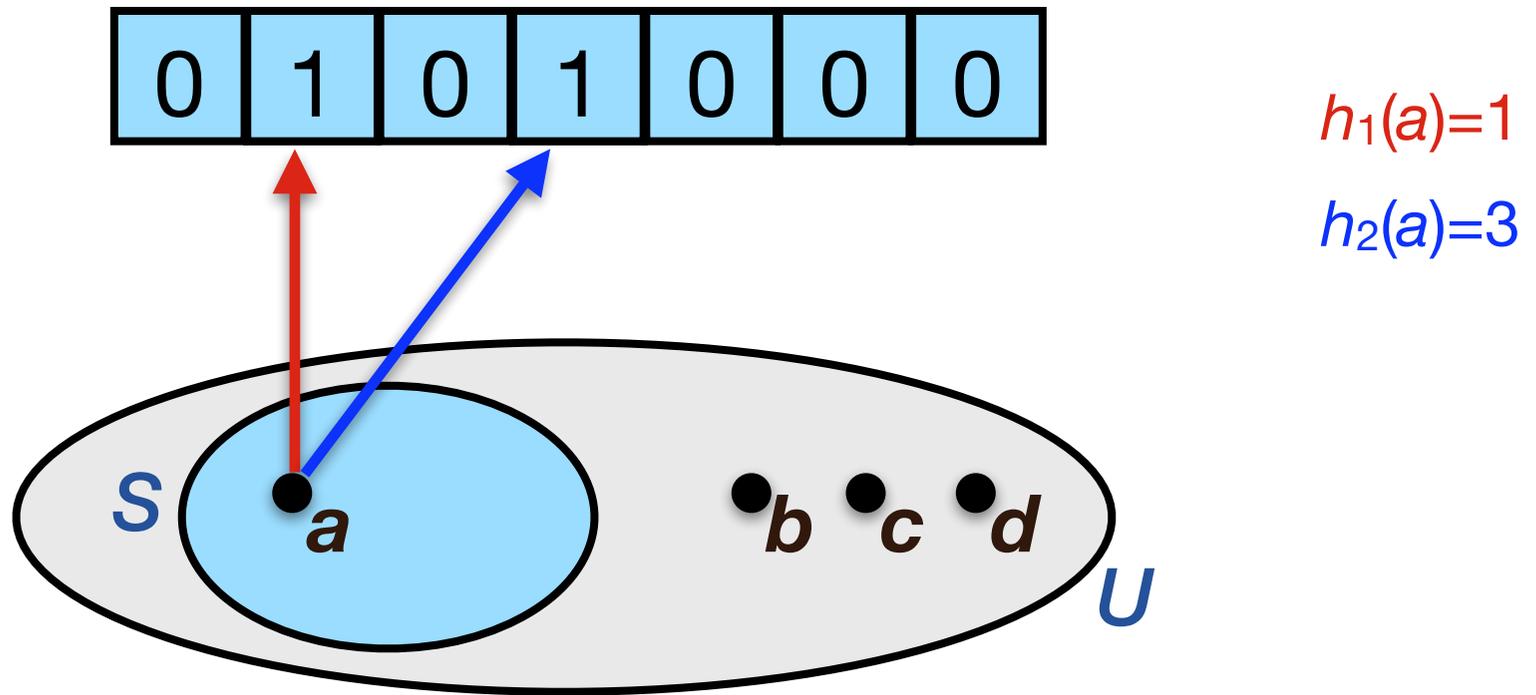
*The most common
filter (by far).*



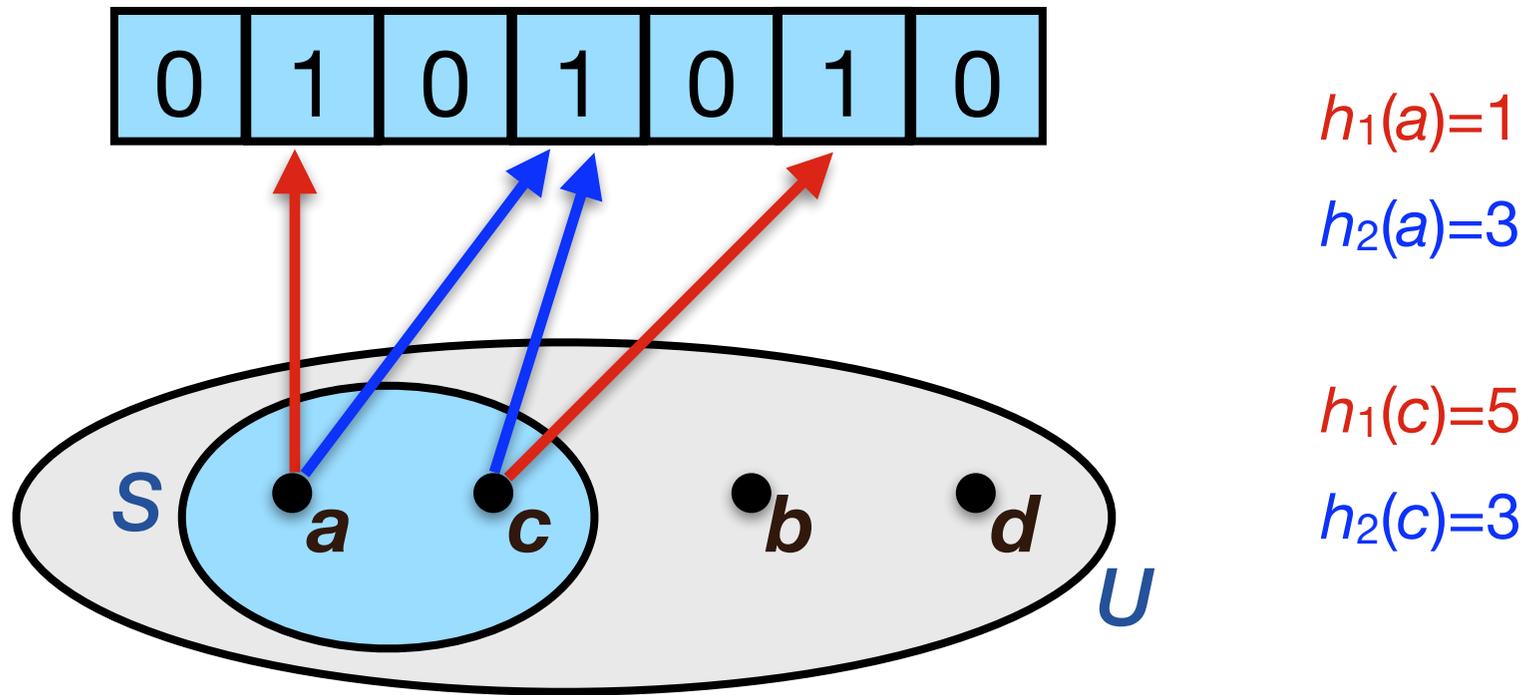
Bloom filter: a bit array + k hash functions. (Here $k=2$.)



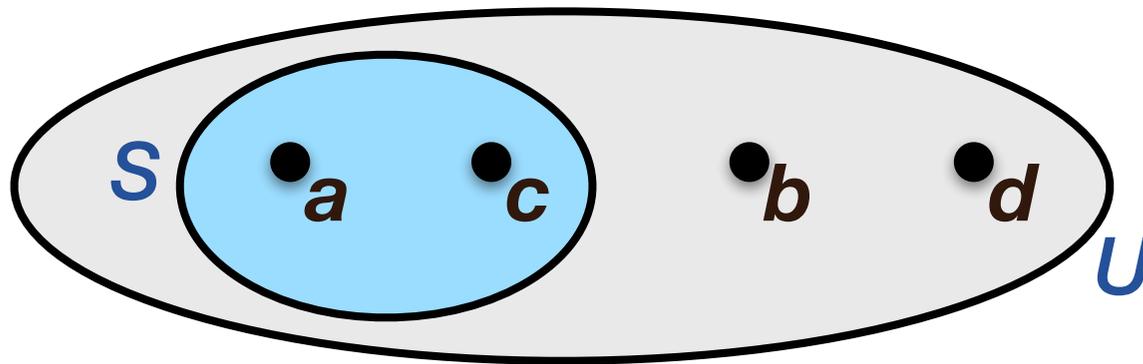
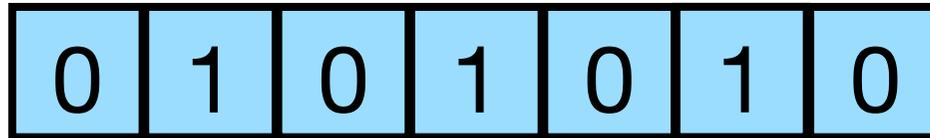
Bloom filter: a bit array + k hash functions. (Here $k=2$.)



Bloom filter: a bit array + k hash functions. (Here $k=2$.)

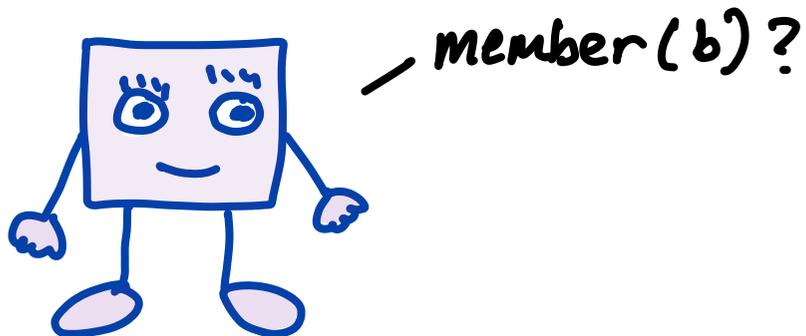
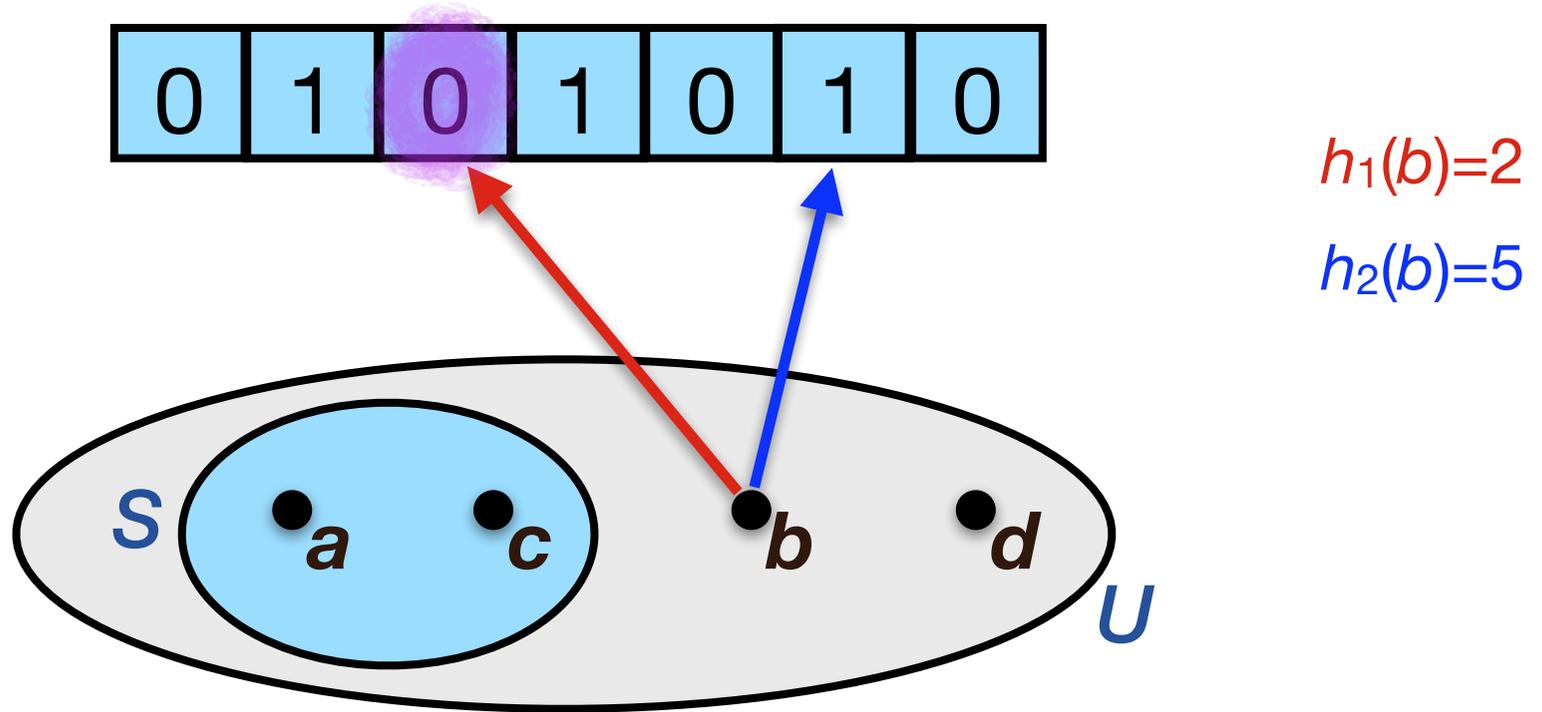


Bloom filter: a bit array + k hash functions. (Here $k=2$.)



Classic Filter: The Bloom Filter [Bloom '70]

Bloom filter: a bit array + k hash functions. (Here $k=2$.)

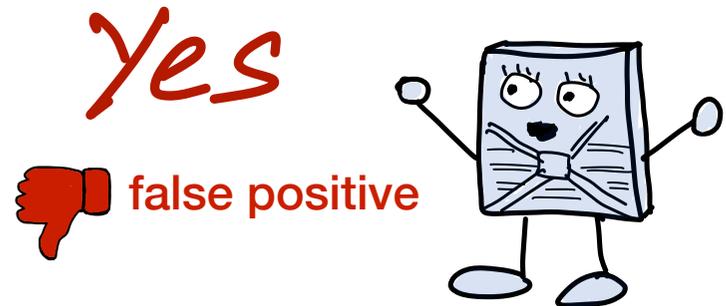
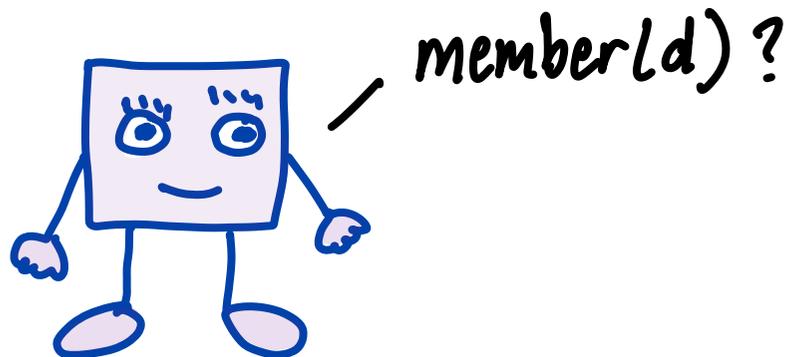
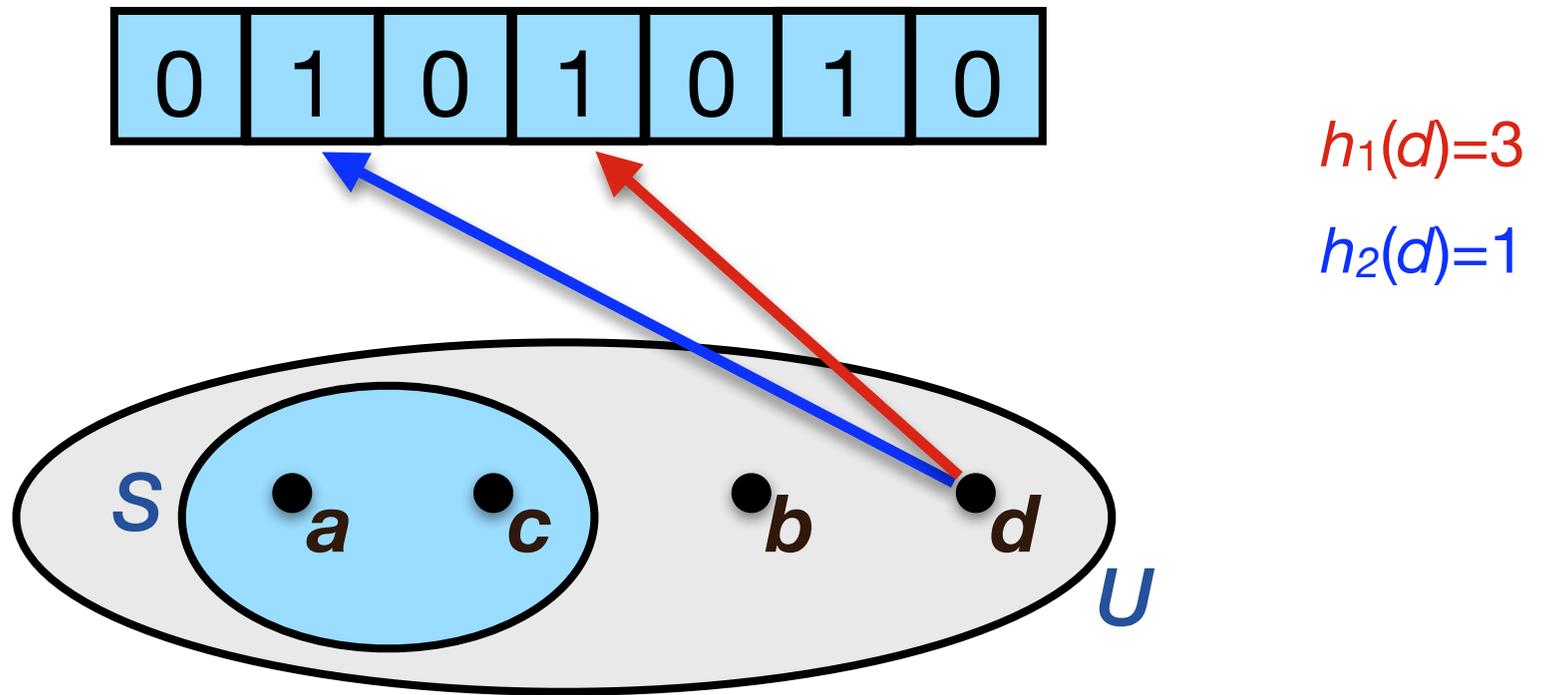


No!

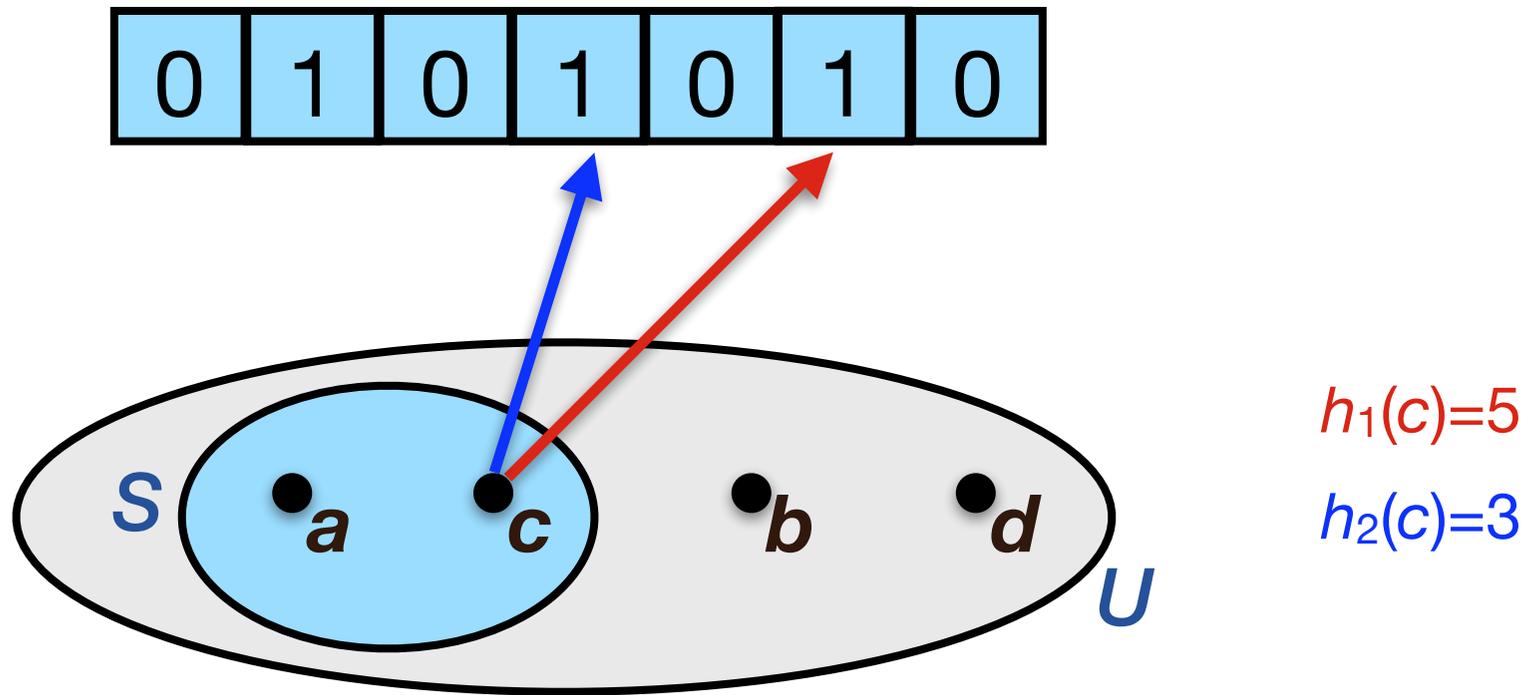


Classic Filter: The Bloom Filter [Bloom '70]

Bloom filter: a bit array + k hash functions. (Here $k=2$.)



Bloom filter: a bit array + k hash functions. (Here $k=2$.)



Bloom filters don't support delete.

Issue: on a delete, which 1s get decremented?

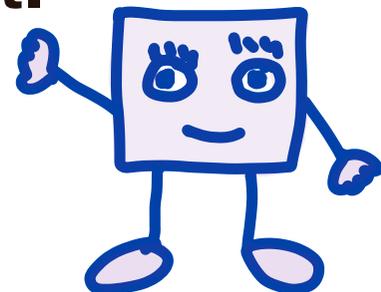
Bloom filter space with false-positive rate ϵ :
 $\approx 1.44 \lg(1/\epsilon)$ bits/element.

Example:

For $\epsilon = 2\%$,
bits/element ≈ 8 .

Common rule of thumb:

Bloom filters take about 1 byte/element.



Bloom filters are ubiquitous

≥4300 citations

Computational biology



Databases



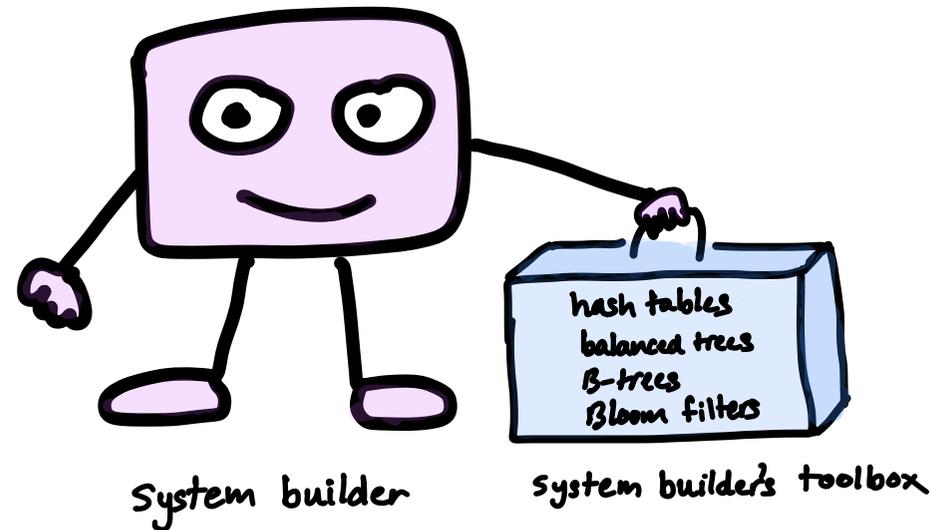
Networking



Storage systems

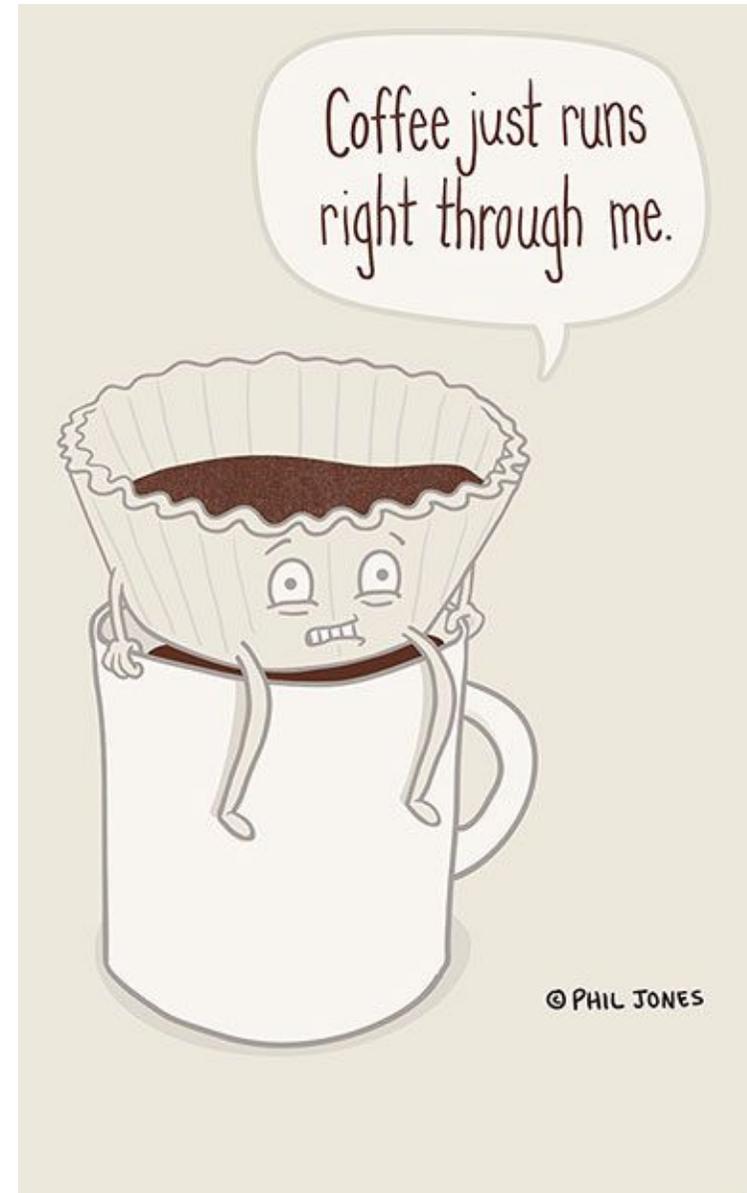


Streaming applications



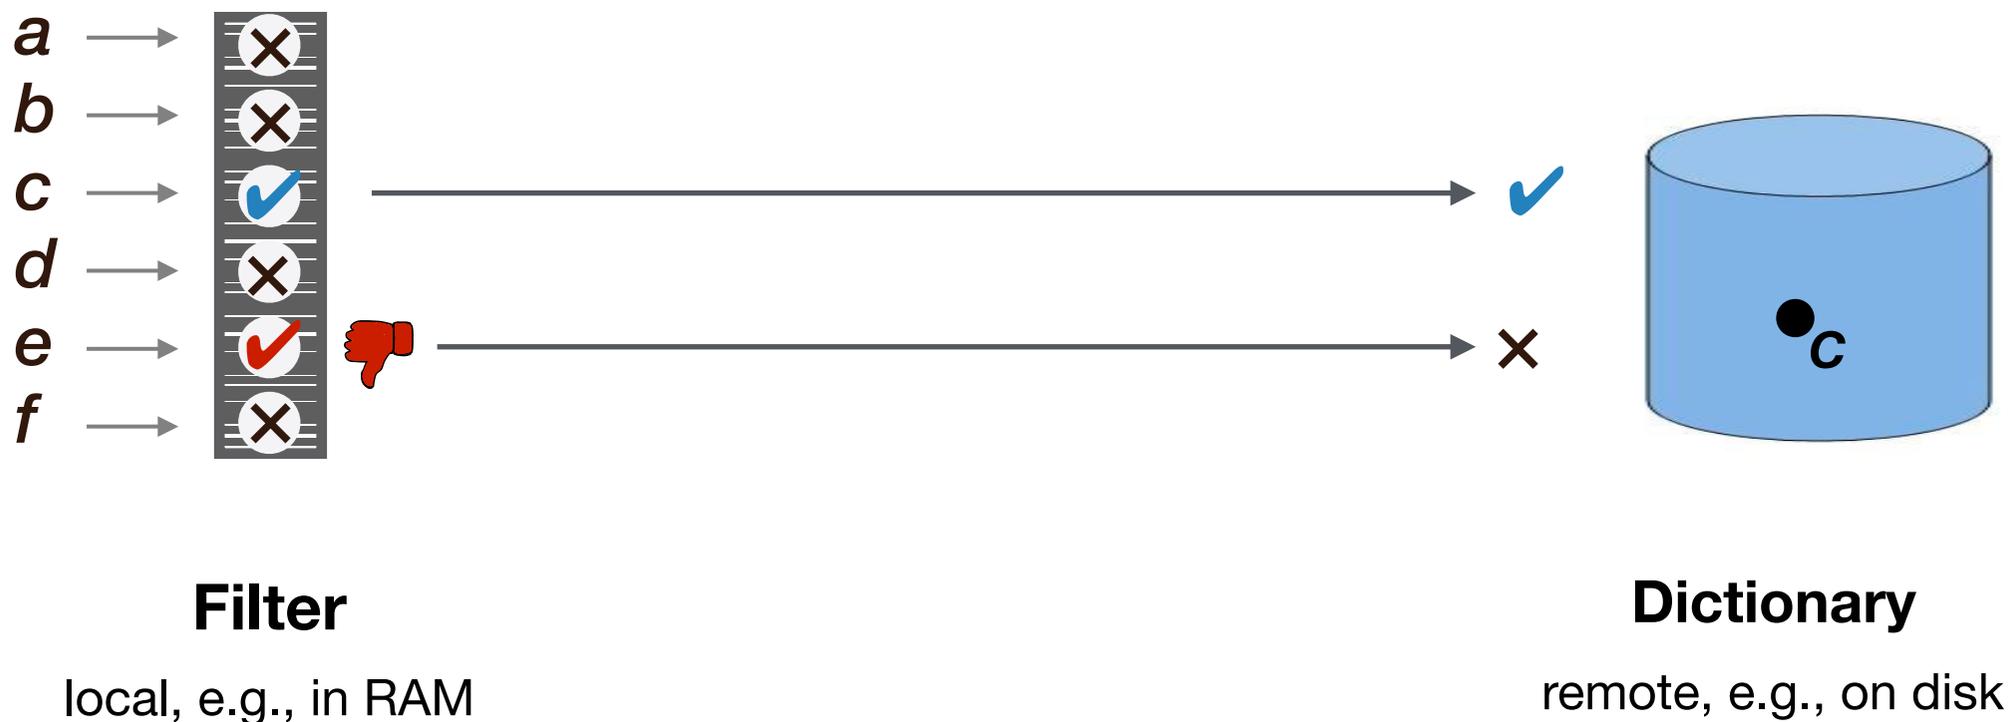
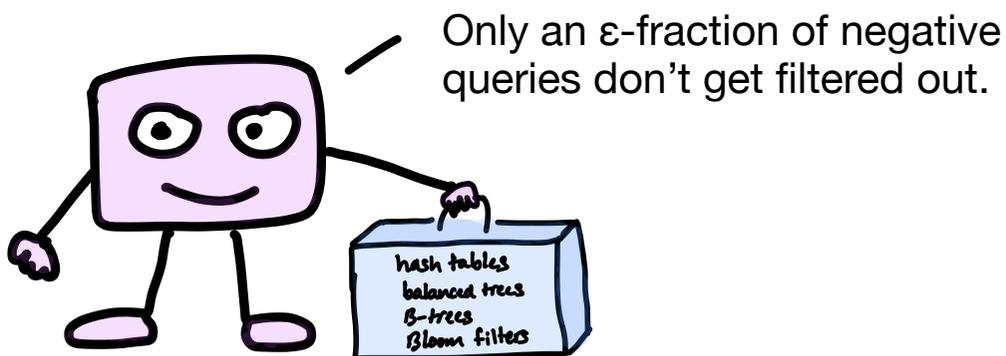
Talk So Far

- Filter data structure
- The Bloom filter [Bloom '70]
- **How filters are used**



Most Common Filter Use

Filter out queries to a large remote dictionary.



Speedup from Filter Use

Workload has P positive and N negative queries.

Dictionaries w/o Bloom Filters	Dictionaries w/ Bloom Filters
$P+N$	$P+\epsilon N$

Remote Accesses of Dictionary

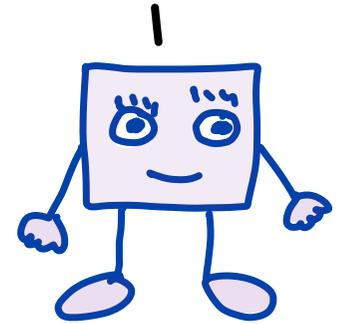
Example: Filters Help Queries in LSM Trees

Log-structured merge tree (LSM)

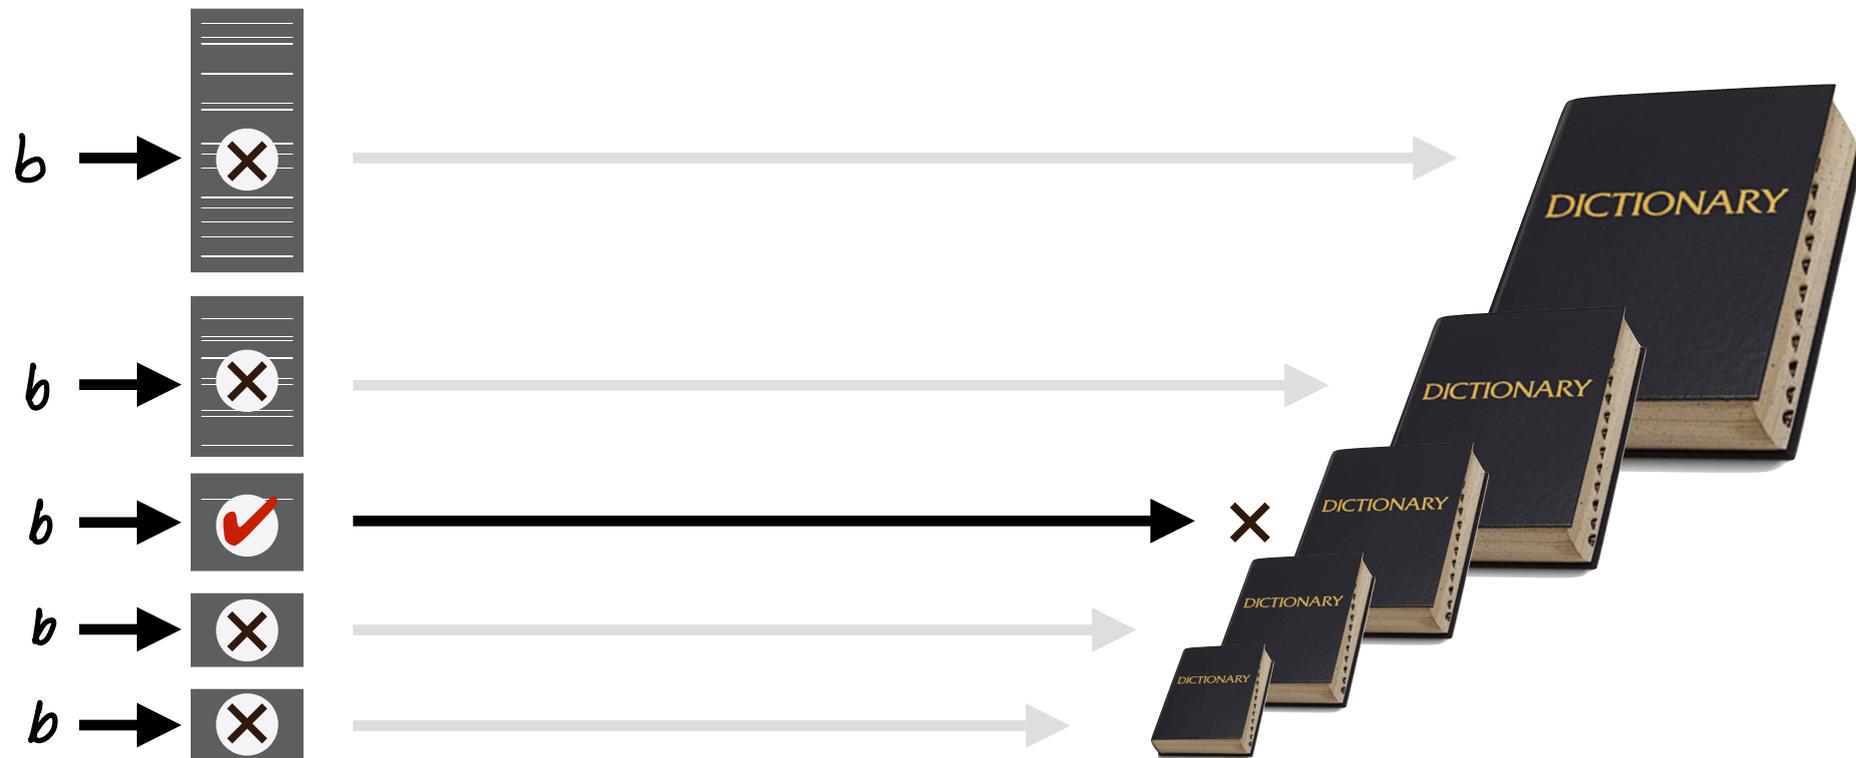
- An LSM tree supports fast inserts by partitioning into independent dictionaries.

[O'Neil, Cheng, Gawlick, O'Neil '96]

member(b)?



Point queries are slow without filters.



Talk So Far

- Filter data structure
- The Bloom filter [Bloom '70]
- How filters are used
- **Time to change your filter**
(the talk title)



Application must work around limited Bloom filter capabilities

Limitations	Work-arounds
No deletes	Rebuild
No resizes	Guess N , and rebuild if wrong
No filter merging nor enumeration of elements	???
No values associated with keys	Combine with other data structure

Bloom filter limitations increase system complexity, waste space, and slow down application performance.

Bloom filters also have suboptimal asymptotics

	Bloom filter	Optimal
Space	$\approx 1.44 n \lg(1/\epsilon)$	$\approx n \lg(1/\epsilon) + O(1)$
CPU cost	$\Omega(\lg(1/\epsilon))$	$O(1)$
Data locality	$\Omega(\lg(1/\epsilon))$ probes	$O(1)$ probes

Bloom filter limitations increase system complexity, waste space, and slow down application performance.

Tons of Research on Extending/Improving/Replacing Bloom Filters

Deletes + counting

[Bonomi, Mitzenmacher, Panigrahy, Singh, Varghese 06],
[Yuan, Miao, Jia, Wang 08],
[Pandey, Bender, Patro, Johnson SIGMOD 17],

Keys

[Chazelle, Kilian, Rubinfeld, Tal 04]

Filters on SSD

[Canim, Mihaila, Bhattacharjee, Lang, Ross 10],
[Debnath, Sengupta, Lilja, Du 11], [Lu, Debnath, Du. 11], [Bender, Farach-Colton, Johnson, Kraner, Kuszmaul, Medjedovic, Montes, Shetty, Spillane, Zadok 12],
[Pandey, Singh, Bender, Berry, Farach-Colton, Johnson, Kroeger, Phillips 20]

Optimizing asymptotics

[Pagh, Pagh, Rao 05],
[Arbitman, Naor, Segev 10],
[Lovett & Porat 10],
[Bender, Farach-Colton, Goswami, Johnson, McCauley, Singh 18]

Engineering

[Bender, Farach-Colton, Johnson, Kraner, Kuszmaul, Medjedovic, Montes, Shetty, Spillane, Zadok 12],
[Fan, Andersen, Kaminsky, Mitzenmacher 14],
[Pandey, Bender, Patro, Johnson SIGMOD 17],
[Pandey, Bender, Johnson, Patro 17],
[Breslow, Jayasena 18],
[Pandey, Conway, Durie, Bender, Martin, Johnson 21]

Adaptivity

[Mitzenmacher, Pontarelli, Reviriego 18]
[Bender, Farach-Colton, Goswami, Johnson, McCauley, Singh 18]
[Bender, Das, Farach-Colton, Mo, Wang 21]

When too much stuff is growing on your filter, it's time to change the filter.

Talk So Far

- Filter data structure
- The Bloom filter [Bloom '70]
- How filters are used
- **Time to change your filter**
(the talk title)

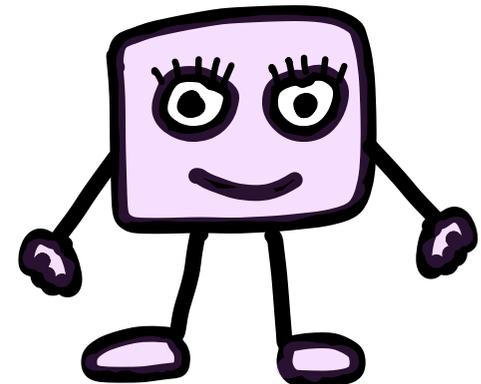


“I see that you change the filter about as often as you change your mind.”

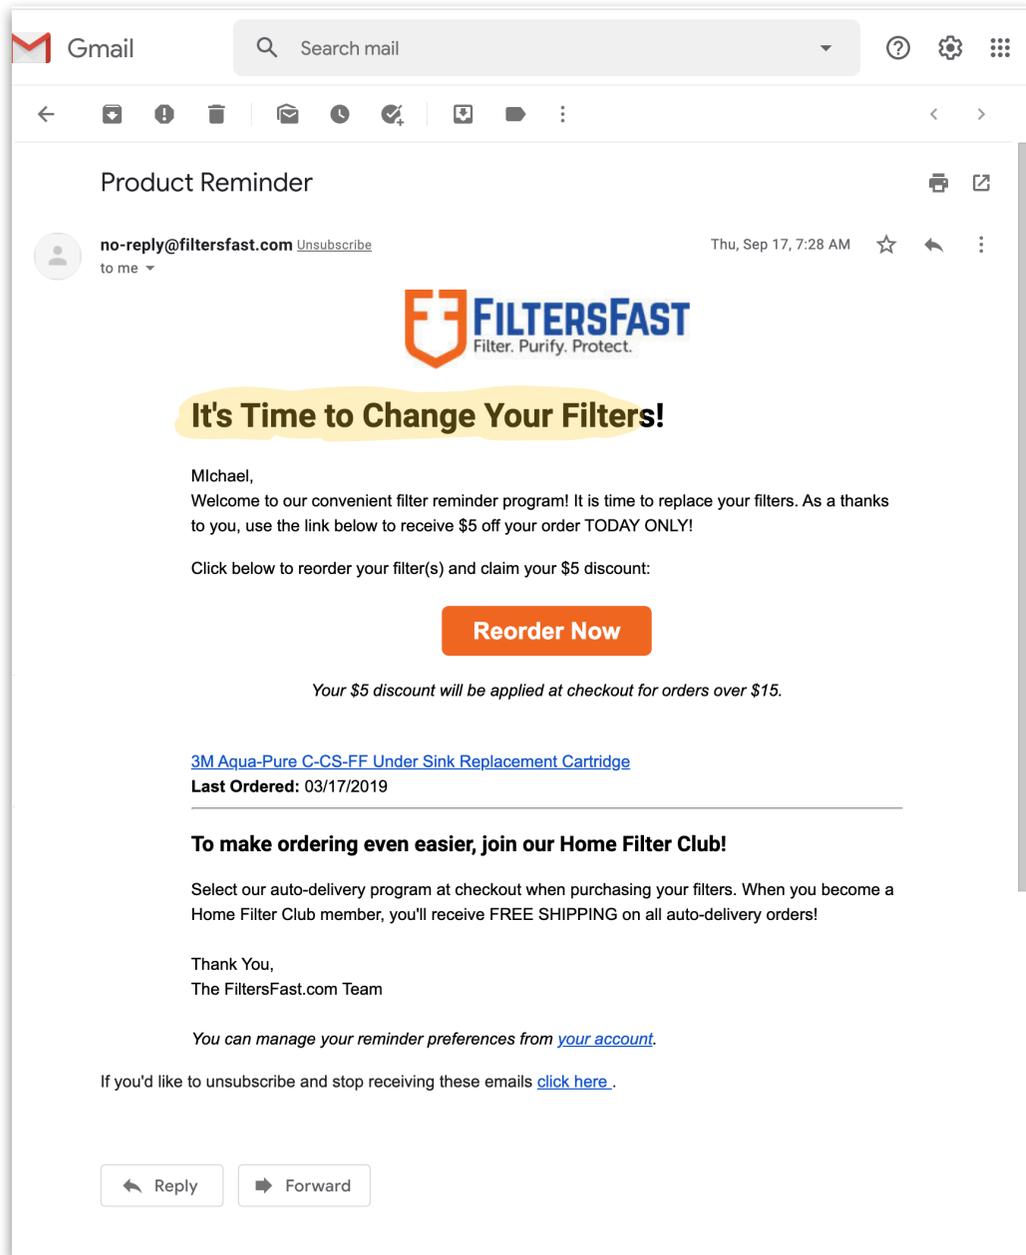
Message of Talk

I couldn't find the
right talk title.

Then I received this email.

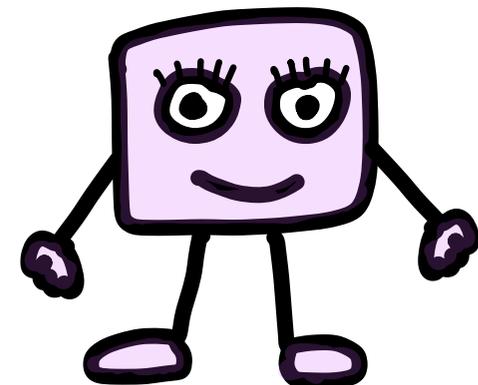


Message of Talk

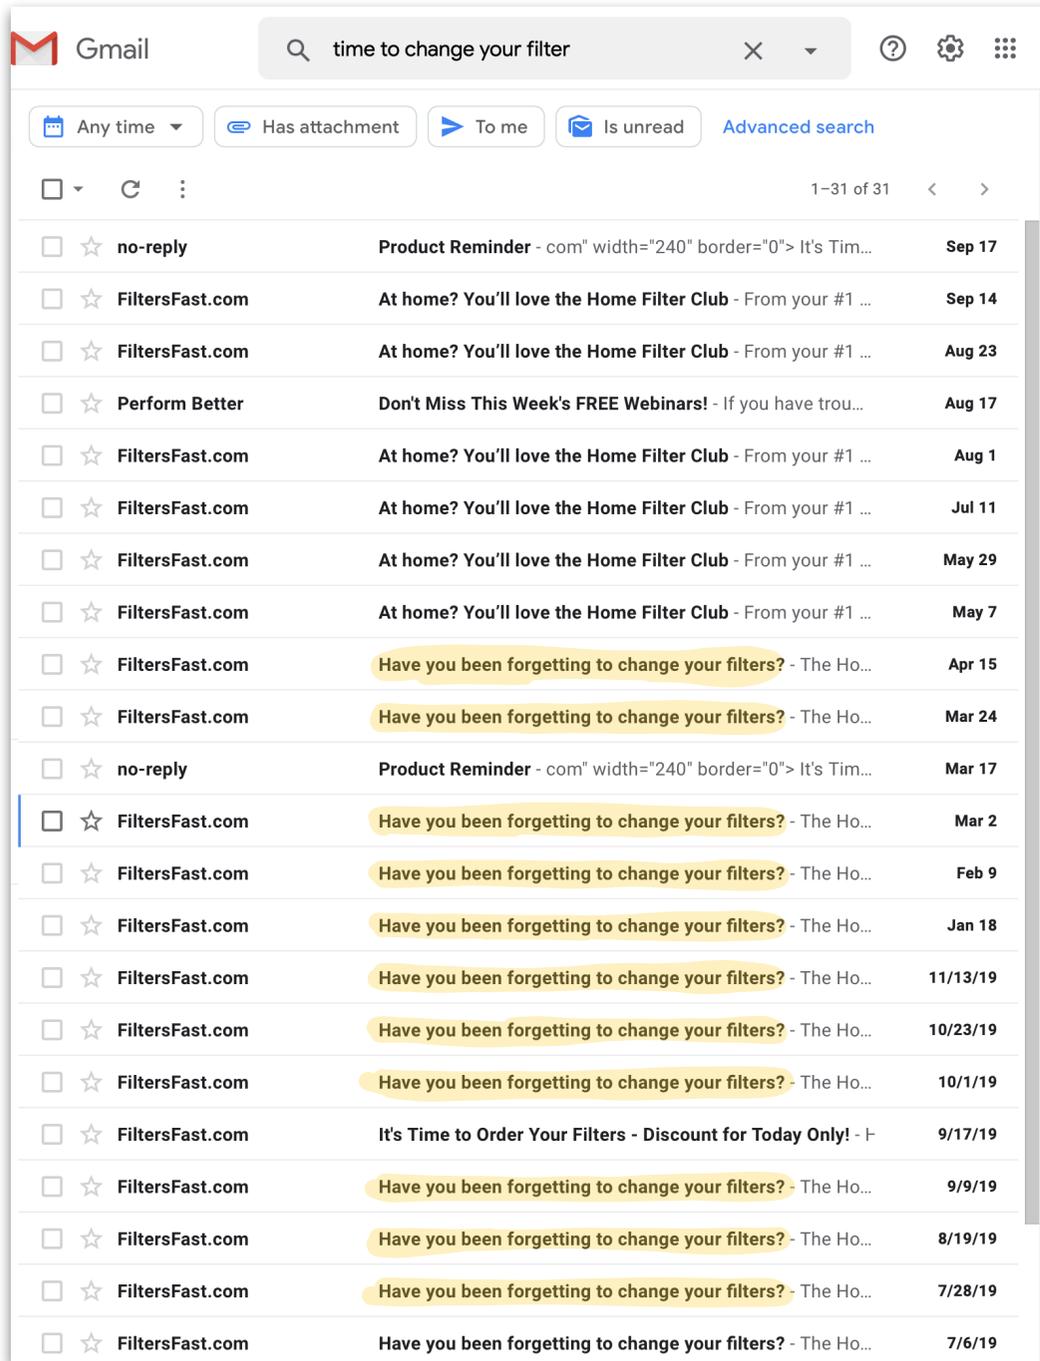


I couldn't find the right talk title.

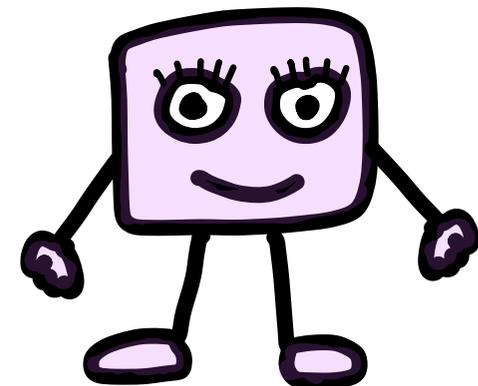
Then I received this email.



Message of Talk



And these emails.

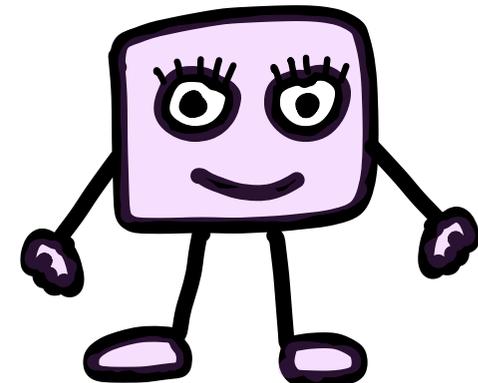


Message of Talk

The screenshot shows a Gmail search results page for the query "time to change your filter". The search filters are set to "Any time", "Has attachment", "To me", and "Is unread". The results list 31 emails, with the first few highlighted in red and the last 15 highlighted in yellow. The highlighted emails are from "FiltersFast.com" and contain the subject "Have you been forgetting to change your filters?".

Sender	Subject	Date
no-reply	Product Reminder - com" width="240" border="0"> It's Tim...	Sep 17
FiltersFast.com	At home? You'll love the Home Filter Club - From your #1 ...	Sep 14
FiltersFast.com	At home? You'll love the Home Filter Club - From your #1 ...	Aug 23
Perform Better	Don't Miss This Week's FREE Webinars! - If you have trou...	Aug 17
FiltersFast.com	At home? You'll love the Home Filter Club - From your #1 ...	Aug 1
FiltersFast.com	At home? You'll love the Home Filter Club - From your #1 ...	Jul 11
FiltersFast.com	At home? You'll love the Home Filter Club - From your #1 ...	May 29
FiltersFast.com	At home? You'll love the Home Filter Club - From your #1 ...	May 7
FiltersFast.com	Have you been forgetting to change your filters? - The Ho...	Apr 15
FiltersFast.com	Have you been forgetting to change your filters? - The Ho...	Mar 24
no-reply	Product Reminder - com" width="240" border="0"> It's Tim...	Mar 17
FiltersFast.com	Have you been forgetting to change your filters? - The Ho...	Mar 2
FiltersFast.com	Have you been forgetting to change your filters? - The Ho...	Feb 9
FiltersFast.com	Have you been forgetting to change your filters? - The Ho...	Jan 18
FiltersFast.com	Have you been forgetting to change your filters? - The Ho...	11/13/19
FiltersFast.com	Have you been forgetting to change your filters? - The Ho...	10/23/19
FiltersFast.com	Have you been forgetting to change your filters? - The Ho...	10/1/19
FiltersFast.com	It's Time to Order Your Filters - Discount for Today Only! - †	9/17/19
FiltersFast.com	Have you been forgetting to change your filters? - The Ho...	9/9/19
FiltersFast.com	Have you been forgetting to change your filters? - The Ho...	8/19/19
FiltersFast.com	Have you been forgetting to change your filters? - The Ho...	7/28/19
FiltersFast.com	Have you been forgetting to change your filters? - The Ho...	7/6/19

And these emails.



This Talk: It's Time to Change Your Filter

- Tutorial-like introduction to filters.
- General techniques for solving filter problems.



fingerprinting



quotienting



collision resolution

This Talk: It's Time to Change Your Filter

- Tutorial-like introduction to filters.
- General techniques for solving filter problems.



fingerprinting



quotienting



collusion
~~collision~~ resolution

This Talk

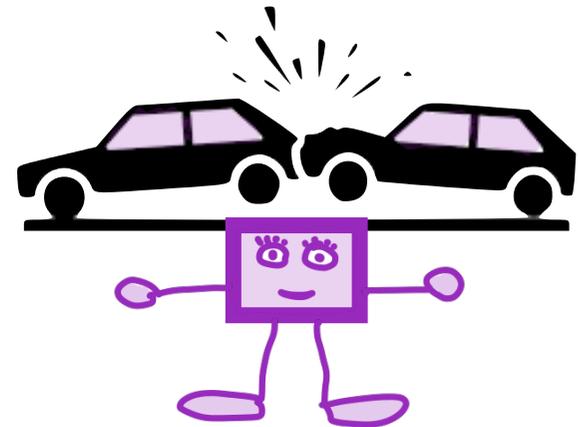
- Tutorial-like introduction to filters.
- General techniques for solving filter problems.



fingerprinting



quotienting



collision resolution

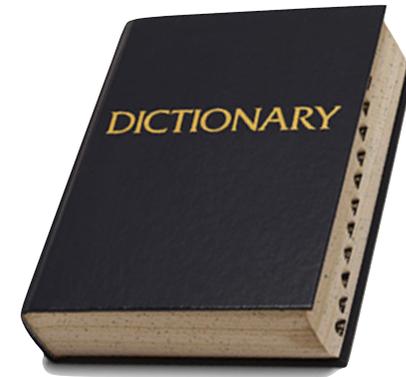
Fingerprinting

[Cleary 84] [Pagh, Pagh, Rao 05], [Arbitman, Naor, Segev 10],
[Bender, Farach-Colton, Johnson, Kraner, Kuszmaul, Medjedovic, Montes,
Shetty, Spillane, Zadok 12], [Fan, Andersen, Kaminsky, Mitzenmacher 14],
[Pandey, Bender, Johnson, Patro 17],
[Bender, Farach-Colton, Goswami, Johnson, McCauley, Singh 18]

Filter F of set S : $\{h(x) \mid x \in S\}$



$$F = \{h(x_1), h(x_2), \dots, h(x_n)\}$$



$$S = \{x_1, x_2, \dots, x_n\}$$

F can be stored more compactly than a Bloom filter.



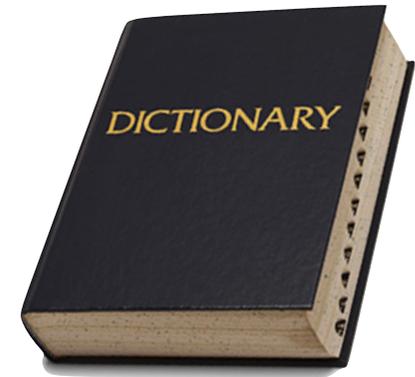
Fingerprinting

[Cleary 84] [Pagh, Pagh, Rao 05], [Arbitman, Naor, Segev 10],
[Bender, Farach-Colton, Johnson, Kraner, Kuszmaul, Medjedovic, Montes,
Shetty, Spillane, Zadok 12], [Fan, Andersen, Kaminsky, Mitzenmacher 14],
[Pandey, Bender, Johnson, Patro 17],
[Bender, Farach-Colton, Goswami, Johnson, McCauley, Singh 18]

Filter F of set S : $\{h(x) \mid x \in S\}$



$$F = \{h(x_1), h(x_2), \dots, h(x_n)\}$$



$$S = \{x_1, x_2, \dots, x_n\}$$

F can be stored more compactly than a Bloom filter.

F is also just dictionary—just a more compact one.



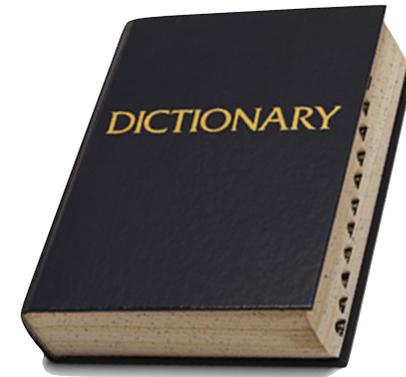
Fingerprinting

[Cleary 84] [Pagh, Pagh, Rao 05], [Arbitman, Naor, Segev 10],
[Bender, Farach-Colton, Johnson, Kraner, Kuszmaul, Medjedovic, Montes,
Shetty, Spillane, Zadok 12], [Fan, Andersen, Kaminsky, Mitzenmacher 14],
[Pandey, Bender, Johnson, Patro 17],
[Bender, Farach-Colton, Goswami, Johnson, McCauley, Singh 18]

Filter F of set S : $\{h(x) \mid x \in S\}$



$$F = \{h(x_1), h(x_2), \dots, h(x_n)\}$$



$$S = \{x_1, x_2, \dots, x_n\}$$

$$\text{member}(x) = \begin{cases} \checkmark & \text{if } h(x) \in F \\ \times & \text{if } h(x) \notin F \end{cases}$$



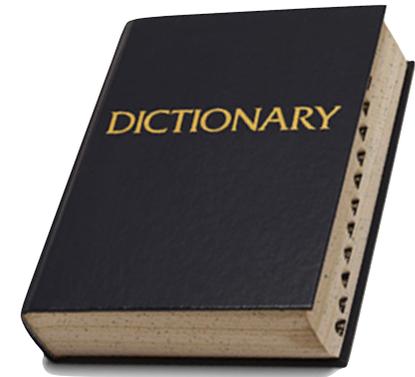
Fingerprinting

[Cleary 84] [Pagh, Pagh, Rao 05], [Arbitman, Naor, Segev 10],
[Bender, Farach-Colton, Johnson, Kraner, Kuszmaul, Medjedovic, Montes,
Shetty, Spillane, Zadok 12], [Fan, Andersen, Kaminsky, Mitzenmacher 14],
[Pandey, Bender, Johnson, Patro 17],
[Bender, Farach-Colton, Goswami, Johnson, McCauley, Singh 18]

Filter F of set S : $\{h(x) \mid x \in S\}$



$$F = \{h(x_1), h(x_2), \dots, h(x_n)\}$$



$$S = \{x_1, x_2, \dots, x_n\}$$

$$\text{member}(x) = \begin{cases} \checkmark & \text{if } h(x) \in F \\ \times & \text{if } h(x) \notin F \end{cases}$$

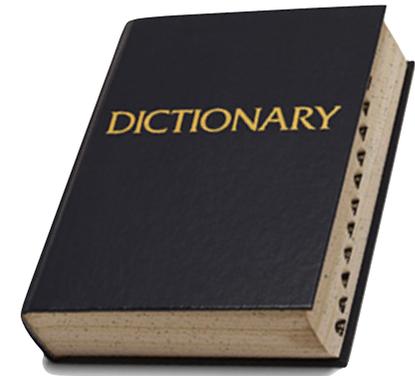
F can be stored more compactly than a Bloom filter.



False-Positive Analysis for Fingerprinting



$$F = \{ h(x_1), h(x_2), \dots, h(x_n) \}$$



$$S = \{ x_1, x_2, \dots, x_n \}$$

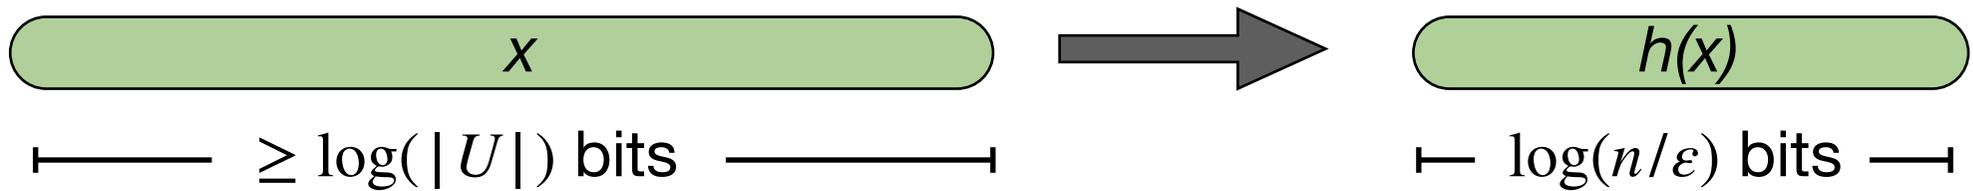
member(y)=✓
if $h(y) \in F$.

y is a false positive
if $\exists i h(y)=h(x_i)$, but $y \notin S$.

Fingerprint collisions: *only* source of false positives.



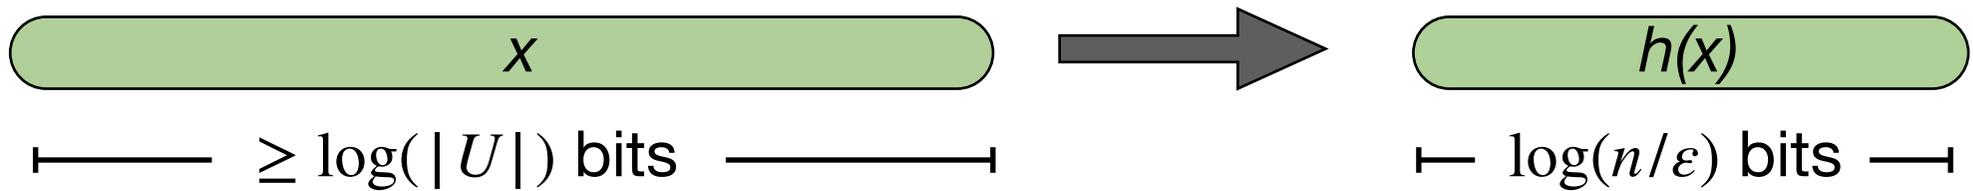
False-Positive Analysis for Fingerprinting



$$\Pr[x \text{ and } y \text{ collide}] = \frac{1}{2^{\log(n/\epsilon)}} = \frac{\epsilon}{n}$$



False-Positive Analysis for Fingerprinting

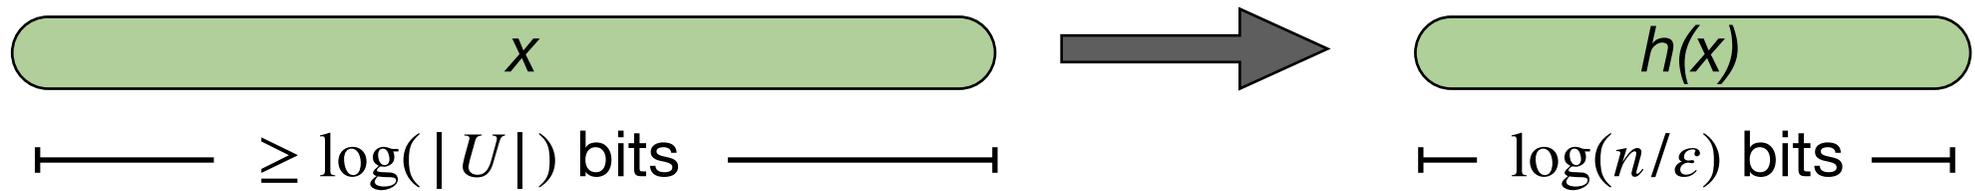


$$\Pr[x \text{ and } y \text{ collide}] = \frac{1}{2^{\log(n/\epsilon)}} = \frac{\epsilon}{n}$$

$$\Pr[y \notin S \text{ is a false positive}] \leq \epsilon$$



False-Positive Analysis for Fingerprinting



$$\Pr[x \text{ and } y \text{ collide}] = \frac{1}{2^{\log(n/\epsilon)}} = \frac{\epsilon}{n}$$

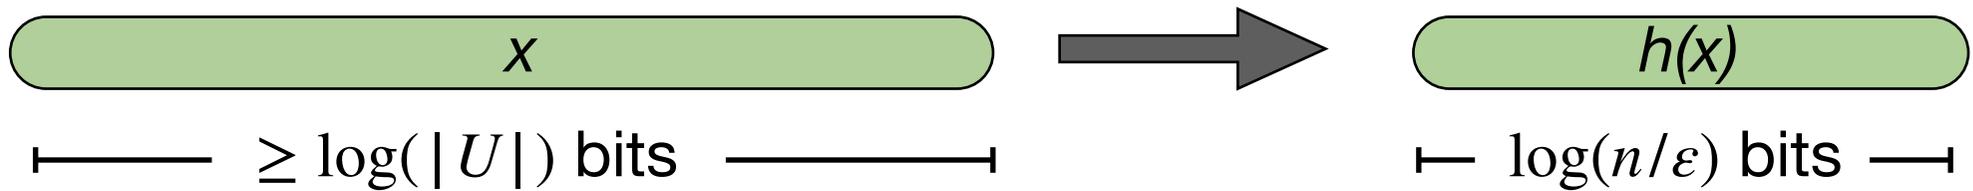
$$\Pr[y \notin S \text{ is a false positive}] \leq \epsilon$$

n fingerprints can be stored compactly:

$$\log(1/\epsilon) + O(1) \text{ bits/element}$$



False-Positive Analysis for Fingerprinting



$$\Pr[x \text{ and } y \text{ collide}] = \frac{1}{2^{\log(n/\epsilon)}} = \frac{\epsilon}{n}$$

$$\Pr[y \notin S \text{ is a false positive}] \leq \epsilon$$

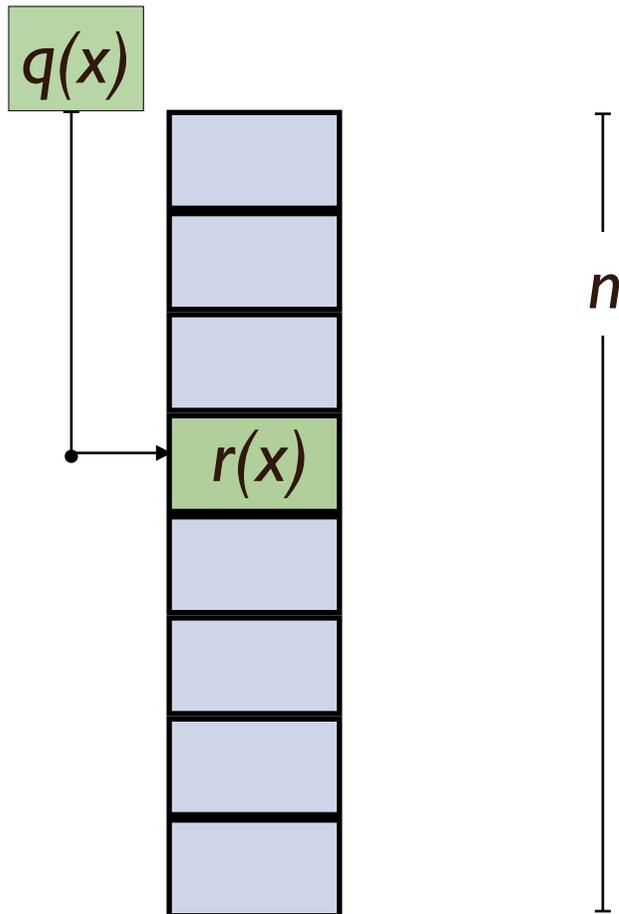
n fingerprints can be stored compactly:

$$\log(1/\epsilon) + O(1) \text{ bits/element}$$



Compact Storage Using Quotienting [Knuth]

Space: $O(\lg(1/\epsilon))$ bits per element



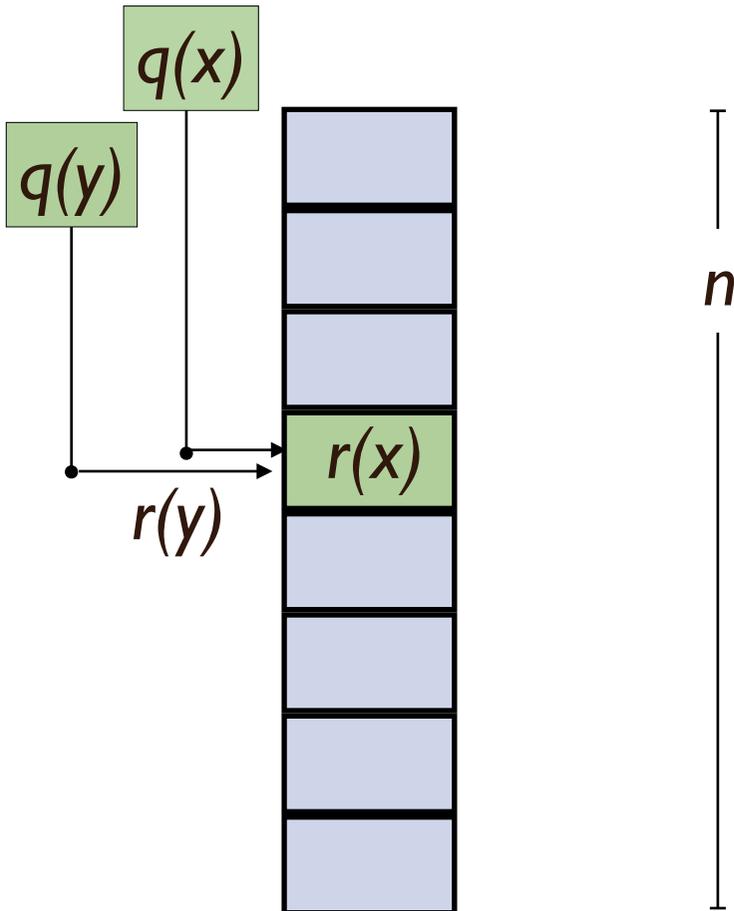
$q(x)$ = location in hash table

$r(x)$ = data stored in hash table



Compact Storage Using Quotienting [Knuth]

Space: $O(\lg(1/\epsilon))$ bits per element



$q(x)$ =location in hash table

$r(x)$ =data stored in hash table

How to deal with collisions in the hash table?

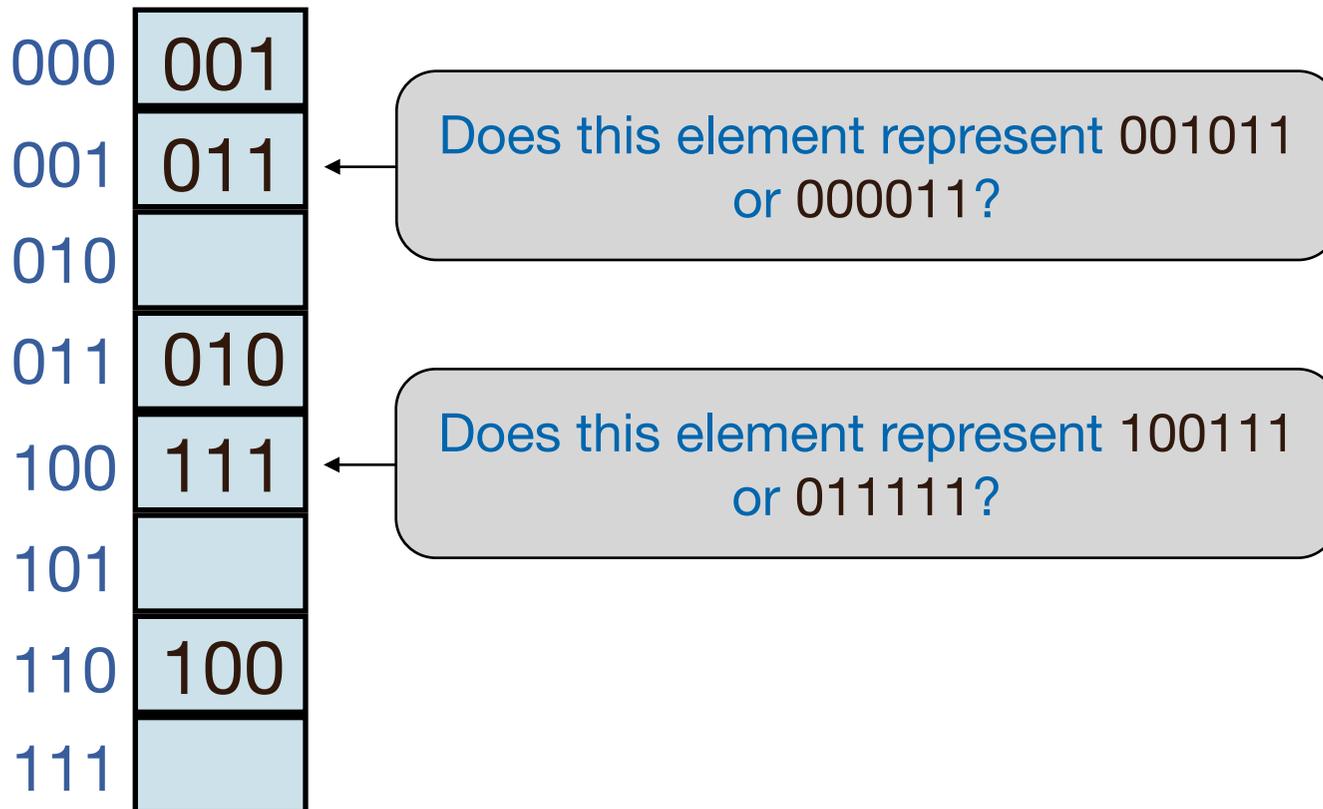
Isn't this a solved problem?

What's wrong with out-of-the-box linear probing?



Hash Collisions in Quotienting

Ex: 6 bit hash. 3 bits for address, 3 for data.

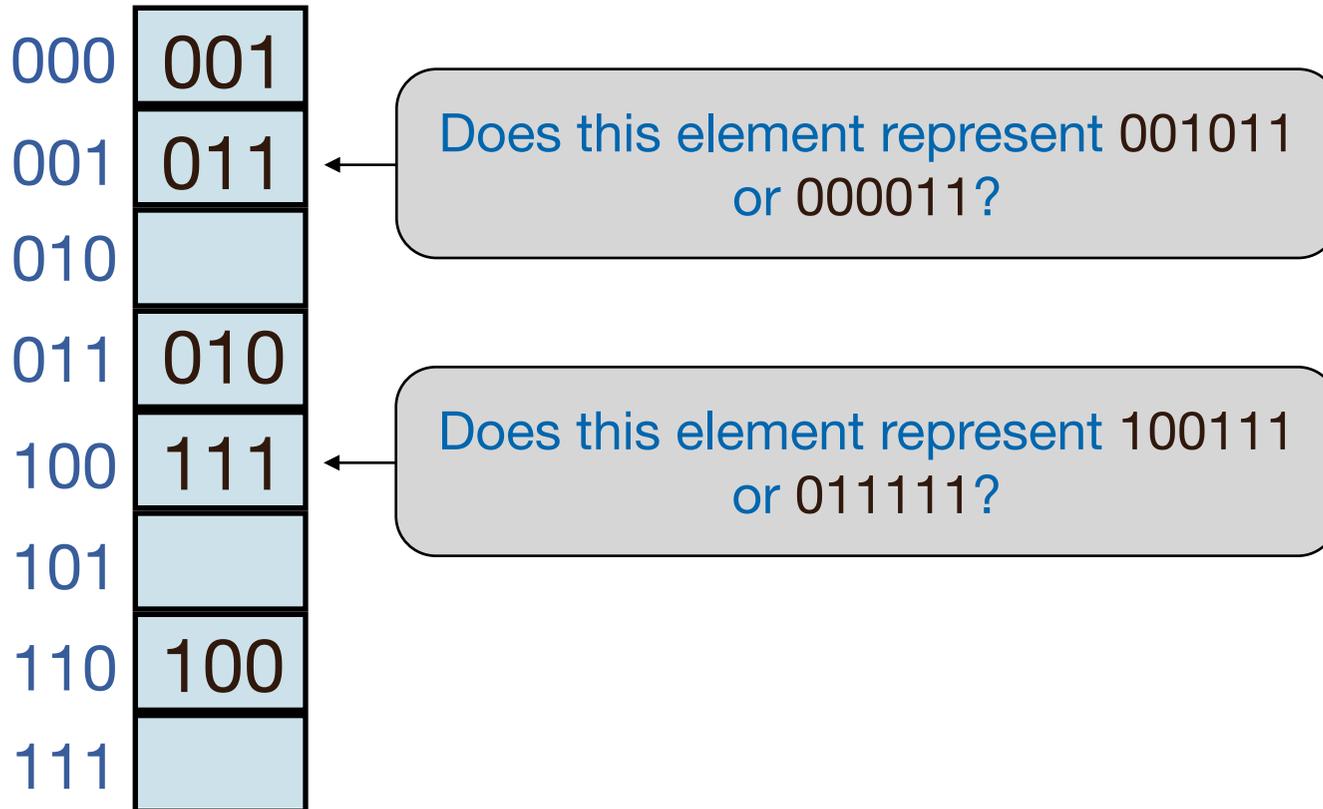


*The hash is stored implicitly based on location.
So how can we change its location?*

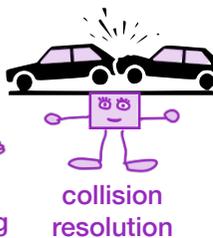


Hash Collisions in Quotienting

Ex: 6 bit hash. 3 bits for address, 3 for data.



*The hash is stored implicitly based on location.
So how can we change its location?*



Talk Structure

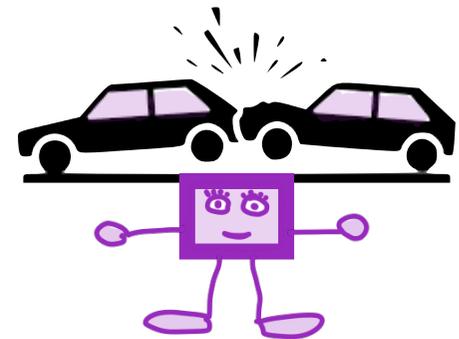
- Filters + how filters are used + Bloom limitations



fingerprinting



quotienting



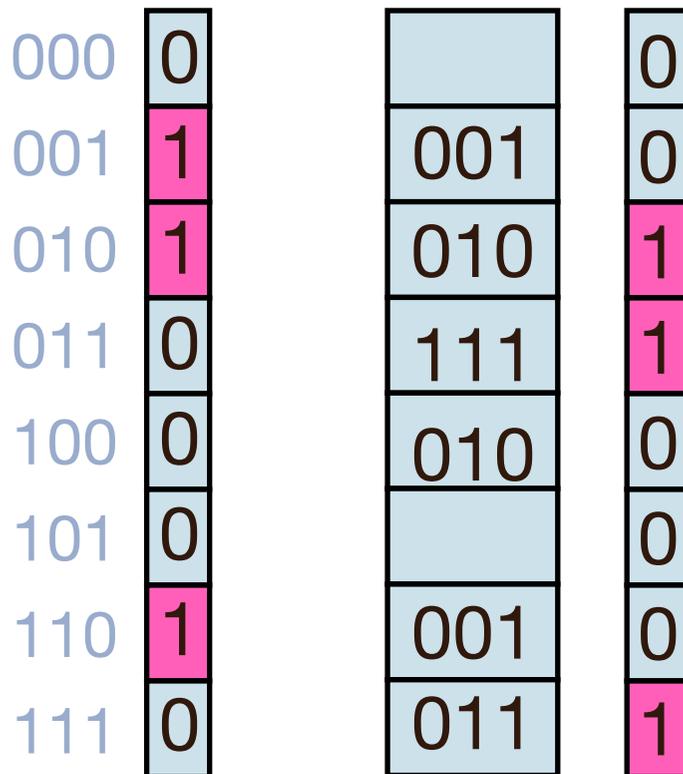
collision resolution

<p>Quotient filters variants</p> <p>[Bender, Farach-Colton, Johnson, Kraner, Kuzmaul, Medjedovic, Montes, Shetty, Spillane, Zadok 12], [Pandey, Bender, Johnson, Patro 17], [Pandey, Bender, Conway, Farach-Colton, Johnson 21]</p>	<p>Cuckoo filter variants</p> <p>[Fan, Andersen, Kaminsky, Mitzenmacher 14], [Breslow, Jayasena 18]</p>	<p>Miscellaneous</p> <p>[Pagh, Pagh, Rao 05], [Arbitman, Naor, Segev 10], [Bender, Farach-Colton, Goswami, Johnson, McCauley, Singh 18]</p>
<p>uses linear proving and Robinhood hashing</p> <p>[Celis, Larson, Munro 85]</p>	<p>Uses Cuckoo hashing</p> <p>[Pagh, Rodler 01]</p>	<p>Balls and bins + more advanced hashing</p>

Quotient Filters

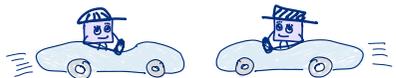
[Bender, Farach-Colton, Johnson, Kraner, Kuzmaul, Medjedovic, Montes, Shetty, Spillane, Zadok VLDB 12]
[Pandey, Bender, Patro, Johnson SIGMOD 17]

2 metadata bits per slot let us recover original location.



1= "something is hashed here"

1= "I'm hashed to the same place as the element before me"

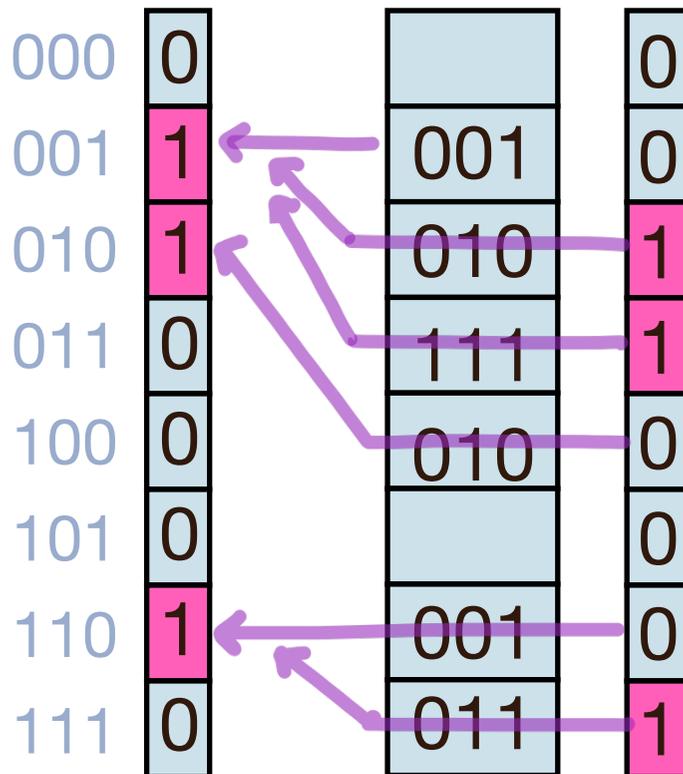


Quotient Filters

[Bender, Farach-Colton, Johnson, Kraner, Kuszmaul, Medjedovic, Montes, Shetty, Spillane, Zadok VLDB 12]
[Pandey, Bender, Patro, Johnson SIGMOD 17]

2 metadata bits per slot let us recover original location.

Recall: no element is stored before its target position.



This is an empty slot.

Remainder is in correct slot.

Remainder is mapped to same location as previous remainder.

Remainder is mapped to same location as previous remainder.

Remainder is mapped to next array position.

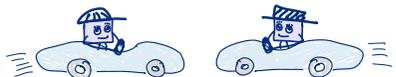
Remainder would map to the next position. But there's none, so it's empty.

Remainder is in correct slot.

Remainder is mapped to same location as previous remainder.

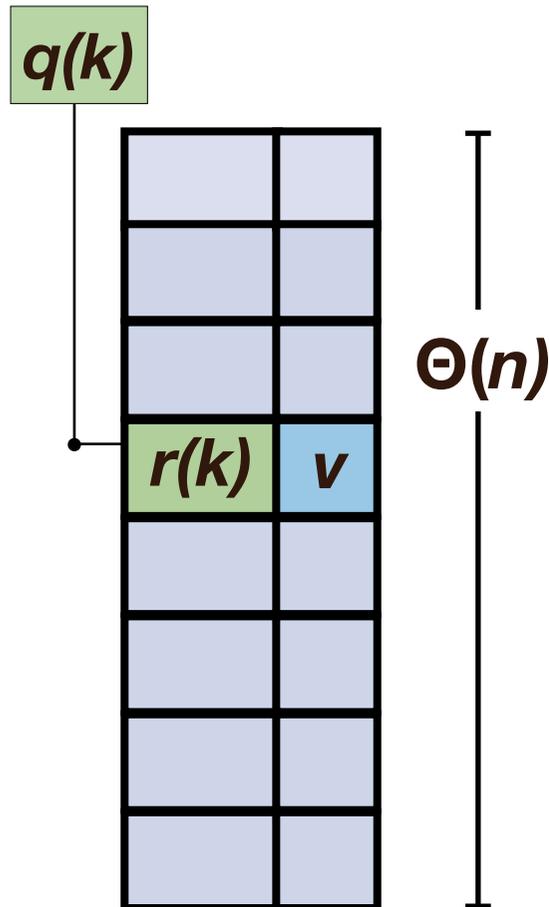
1= "something is hashed here"

1= "I'm hashed to the same place as the element before me"



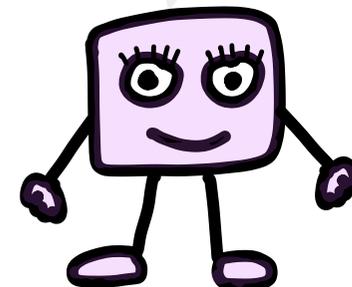
Quotient Filter Capabilities

[Bender, Farach-Colton, Johnson, Kraner, Kuszmaul, Medjedovic, Montes, Shetty, Spillane, Zadok VLDB 12]
[Pandey, Bender, Patro, Johnson SIGMOD 17]



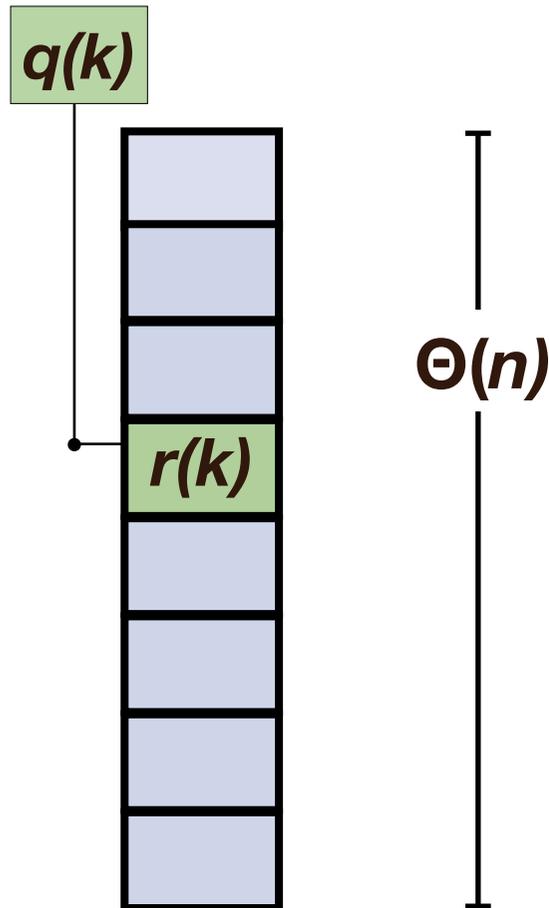
Bloom limitations
No deletes or counting
No resizes
No element enumeration or merging of filters
Keys have no values

QF Capabilities
counting and deletes
resizing
element enumeration + filter merging
values allowed



Quotient Filter Capabilities

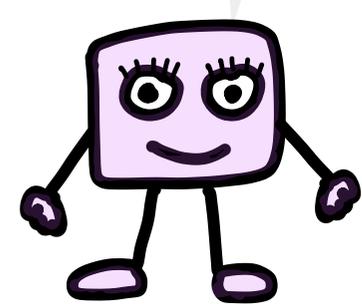
[Bender, Farach-Colton, Johnson, Kraner, Kuszmaul, Medjedovic, Montes, Shetty, Spillane, Zadok VLDB 12]
 [Pandey, Bender, Patro, Johnson SIGMOD 17]



	Bloom performance
Space	$\approx 1.44 n \lg(1/\epsilon)$
CPU cost	$\Omega(\lg(1/\epsilon))$
Data locality	$\Omega(\lg(1/\epsilon))$ probes

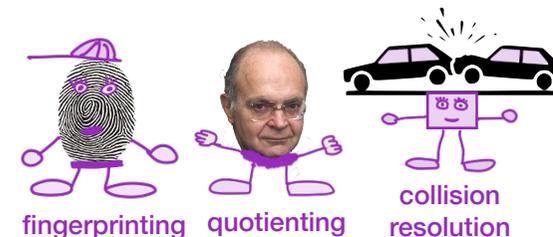
QF performance
 $(1 + \text{tiny})n$
 $\approx 1.44n(\lg(1/\epsilon) + 2)$
 $O(1)$ expected
 1 probe + scan

QF has practical + theoretical performance advantages.



Reminder of General Approach

- **Fingerprinting:** key $k \in S \rightarrow h(k) \in F$.
 - ▶ False-positives only come from fingerprint collisions.
- **Quotienting:** Store fingerprints in a hash table implicitly.
 - ▶ $\lg n$ bits of fingerprint depend on hash location.
- **Collision resolution:**
 - ▶ E.g., using linear probing/Robinhood hashing.
 - ▶ Use metadata bits to recover each $h(k)$.
- **Designs alternatives:** cuckoo, Mortan, broom
 - ▶ Use a different hash table but the same general approach



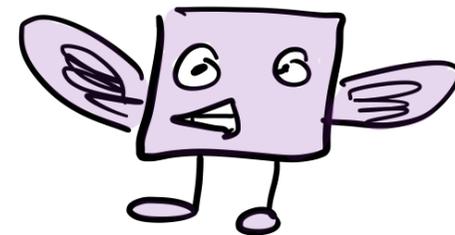
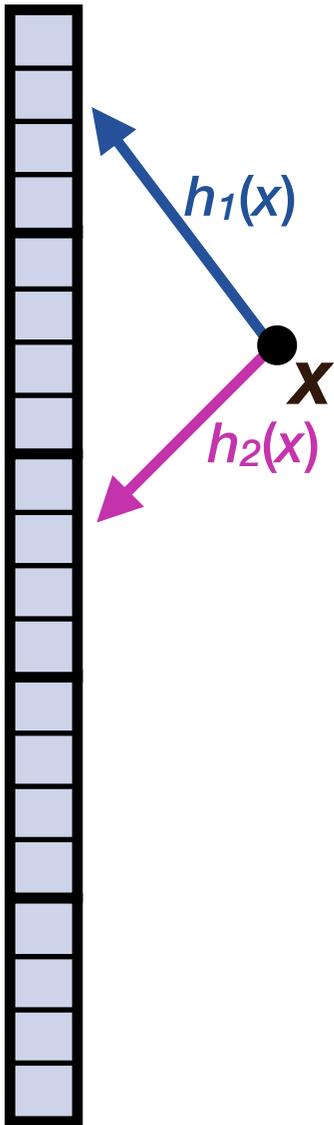
Cuckoo Hashing → Cuckoo Filters

[Pagh, Rodler 01]

[Fan, Andersen, Kaminsky, Mitzenmacher 14] [Breslow, Jayasena 18]

Cuckoo hash table has 2 hash functions h_1 and h_2 .

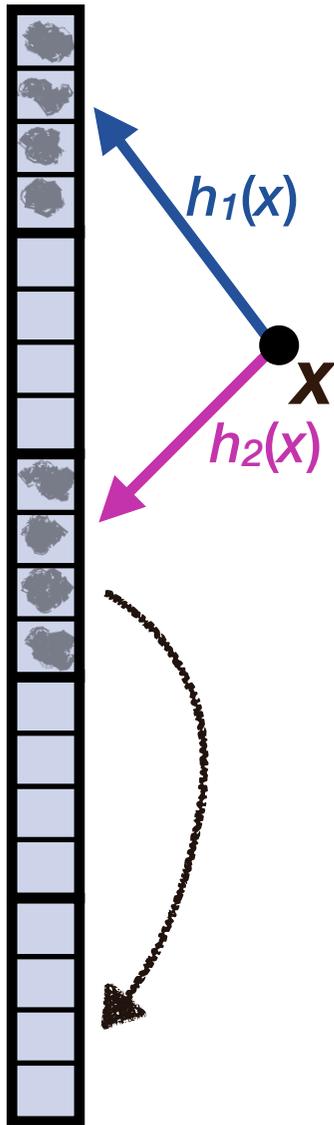
Each hash bucket has 4 slots.



Cuckoo Hashing → Cuckoo Filters

[Pagh, Rodler 01]

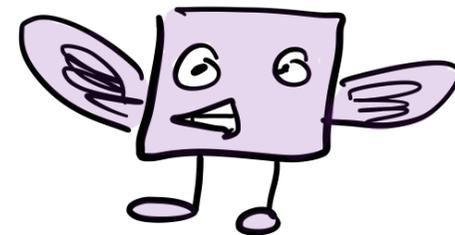
[Fan, Andersen, Kaminsky, Mitzenmacher 14] [Breslow, Jayasena 18]



Cuckoo hash table has 2 hash functions h_1 and h_2 .

Each hash bucket has 4 slots.

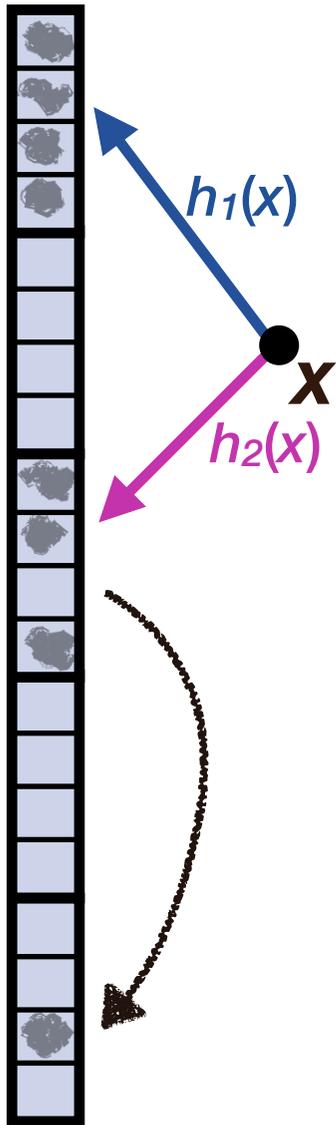
If there's no space in any of the 8 slots:
kick out an element, and
move it to the alternative location
(which may cause other kicks).



Cuckoo Hashing → Cuckoo Filters

[Pagh, Rodler 01]

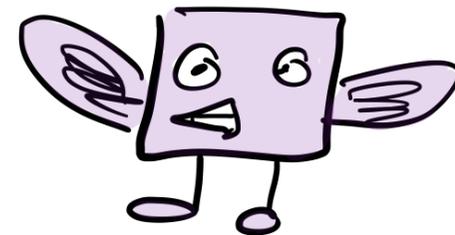
[Fan, Andersen, Kaminsky, Mitzenmacher 14] [Breslow, Jayasena 18]



Cuckoo hash table has 2 hash functions h_1 and h_2 .

Each hash bucket has 4 slots.

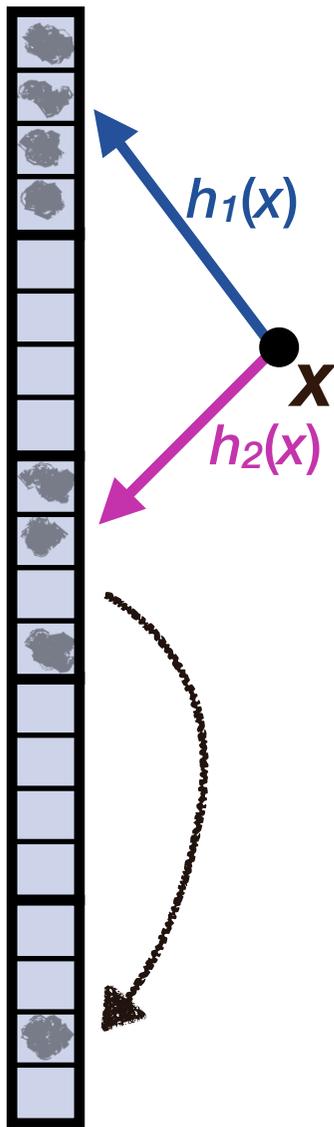
If there's no space in any of the 8 slots:
kick out an element, and
move it to the alternative location
(which may cause other kicks).



Obstacles for Cuckoo Filters

[Pagh, Rodler 01]

[Fan, Andersen, Kaminsky, Mitzenmacher 14] [Breslow, Jayasena 18]



Q: If $f(x)$ is kicked, how to find an alternative when don't store x ?

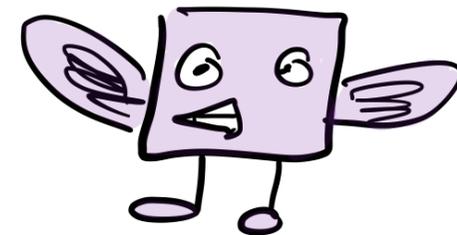
A: We give up on independent hash functions. The alternate location only depends on $r(x)$.

We also give up on asymptotic correctness.

Amazingly, it works for (practical) n not too large.

Cuckoo hashing seemingly doesn't have metadata bits.

But because there are 4 slots per cell, 2 more fingerprint bits are stored explicitly.



Quotient Filter and Cuckoo Filter Comparison

Quotient filter	Cuckoo filter
good space	good space
very good locality	ok locality
some degradation at high load factors	degradation at high load factors
good searches	very good searches
fast inserts	fast inserts
supports counting, multisets, values, deletes	
complicated code	code is easier
good for all n .	pre-asymptotic guarantees. fails w.h.p. large enough n

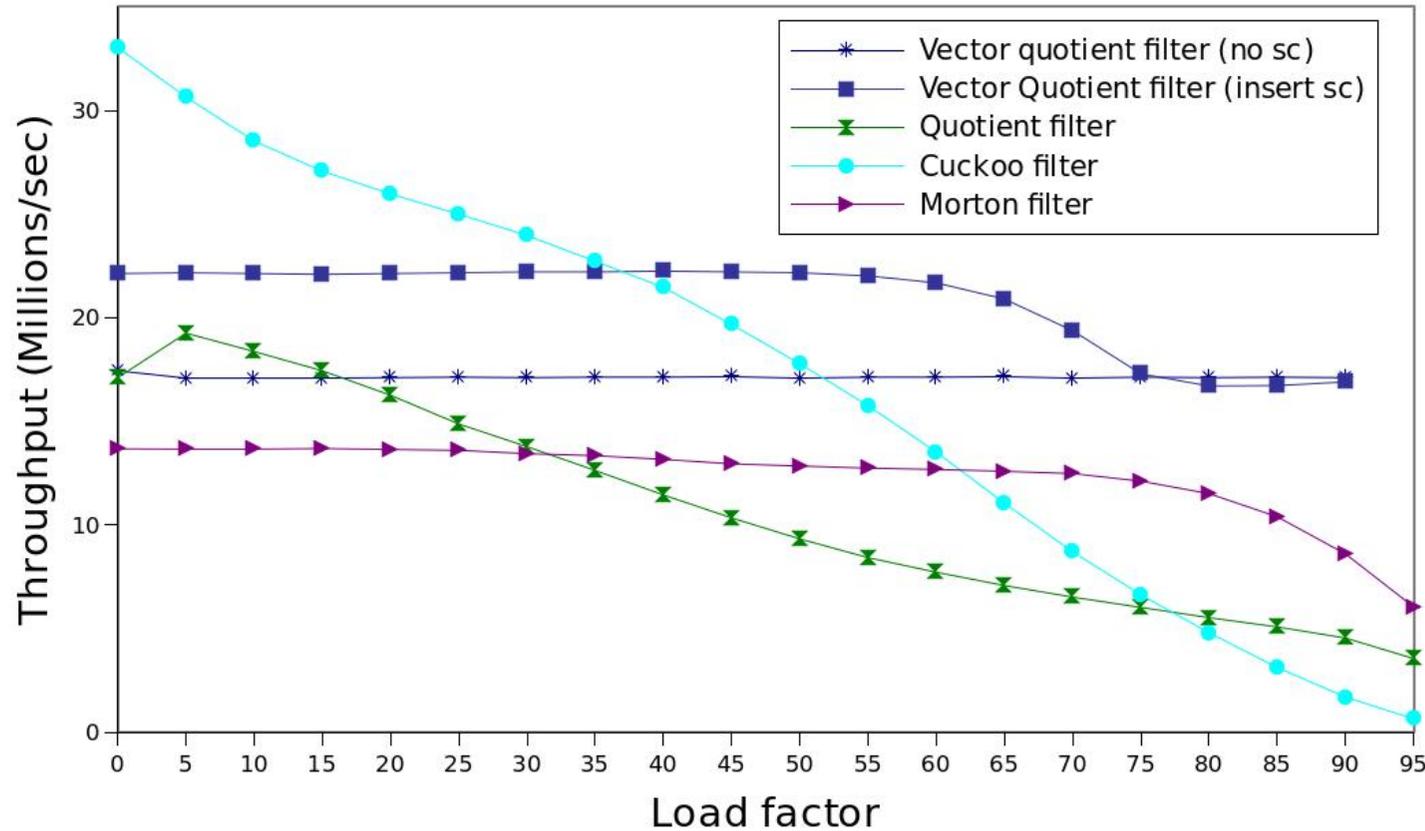
Theorem: There is an optimal filter with

- ▶ **Space:** $(1 + o(1)) n \log(1/\varepsilon) + O(n)$
- ▶ **Error rate:** $\leq \varepsilon$
- ▶ **Operations:** $O(1)$ insert, delete, query.

- [Pagh, Pagh, Rao 05]
- [Arbitman, Naor, and Segev 10]
- [Lovett & Porat 10]
- [Bender, Farach-Colton, Goswami, Johnson, McCauley, Singh 18] adaptive & worst case

Empirical Performance of a Vector Quotient Filter

[Pandey, Conway, Durie, Bender, Farach-Colton, Johnson 21

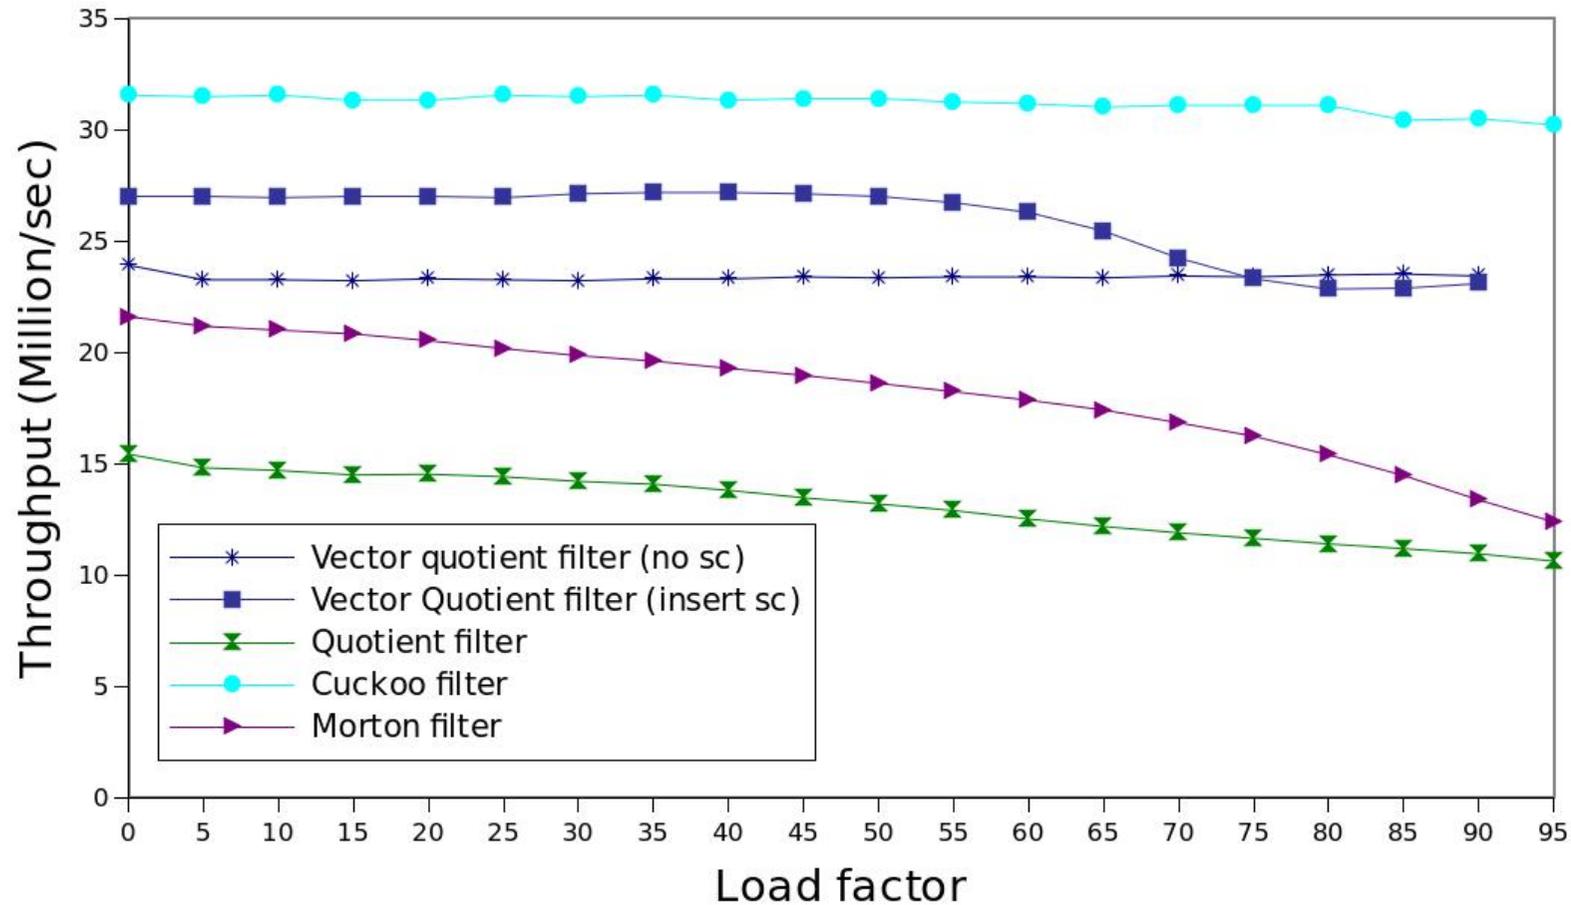


(a) Insertion (Higher is better.)

Quotient filter plus minimum of 2 choice retains good performance even when almost full.

Empirical Performance of a Vector Quotient Filter

[Pandey, Conway, Durie, Bender, Farach-Colton, Johnson 21



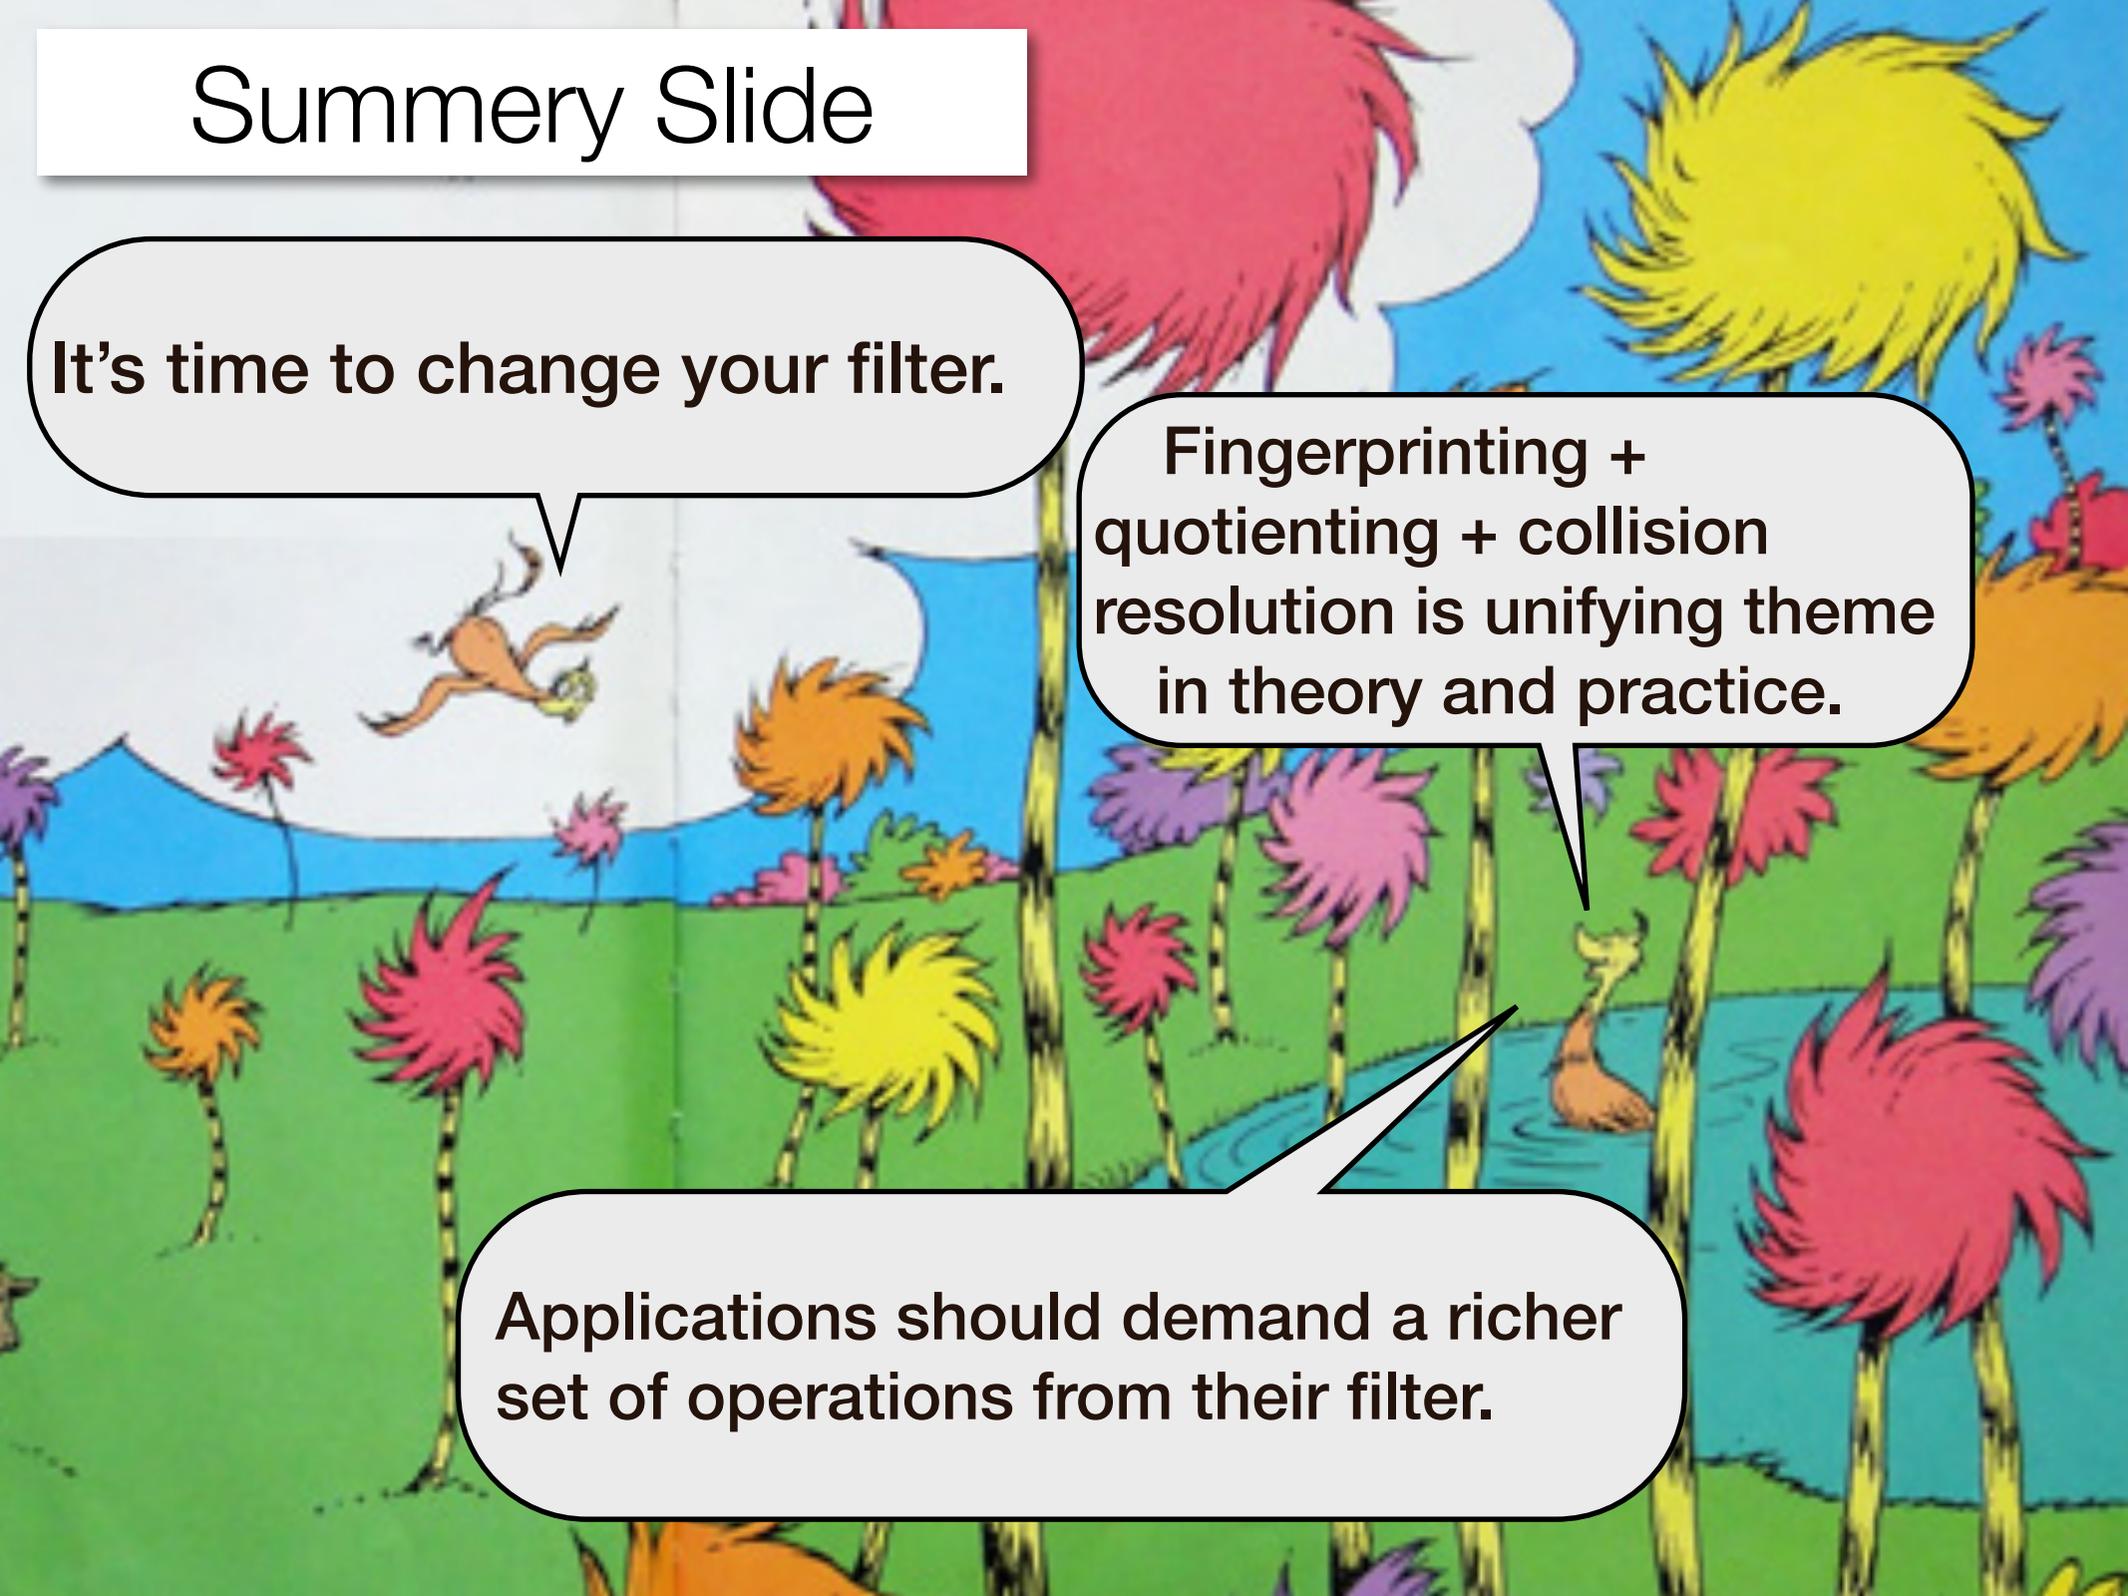
(c) Successful lookup (Higher is better.)

Cuckoo still kicks butt on queries.

Summery Slide



Summery Slide



It's time to change your filter.

Fingerprinting +
quotienting + collision
resolution is unifying theme
in theory and practice.

Applications should demand a richer
set of operations from their filter.