

Data Structures and Algorithms

Bloom Filters

CS 225
Brad Solomon

November 19, 2025



Department of Computer Science

Learning Objectives

Review when you would prefer different data structures

Build a conceptual understanding of a bloom filter

Review probabilistic data structures and one-sided error

Formalize the math behind the bloom filter

Running Times

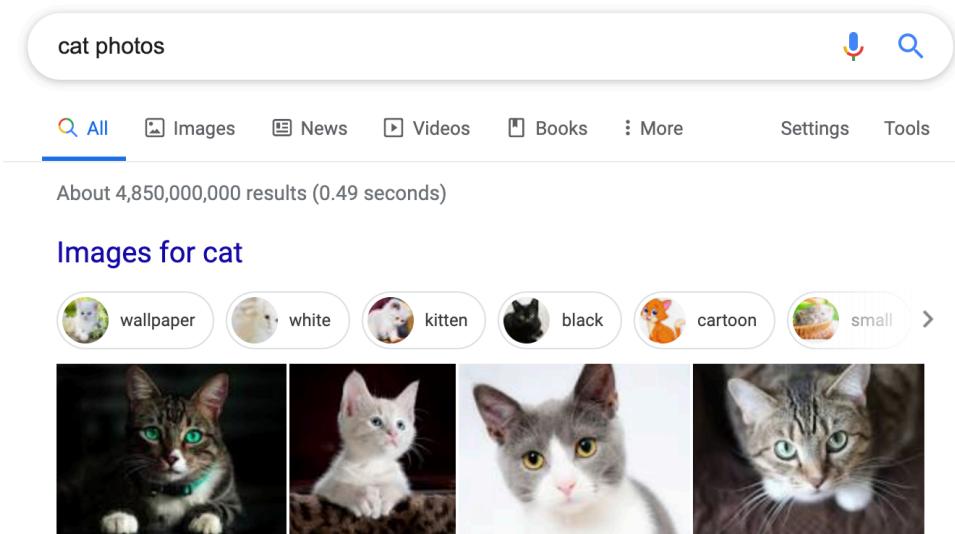


	Hash Table	AVL	Linked List
Find	Expectation*: $O(1)^{***}$ Worst Case: $O(n)$	$O(\log n)$	$O(n)$
Insert	Expectation*: $O(1)^{***}$ Worst Case: $O(n)$	$O(\log n)$	$O(1)$
Storage Space	$O(n)$	$O(n)$	$O(n)$

Memory-Constrained Data Structures

What method would you use to build a search index on a collection of objects *in a memory-constrained environment?*

Constrained by Big Data (Large N)



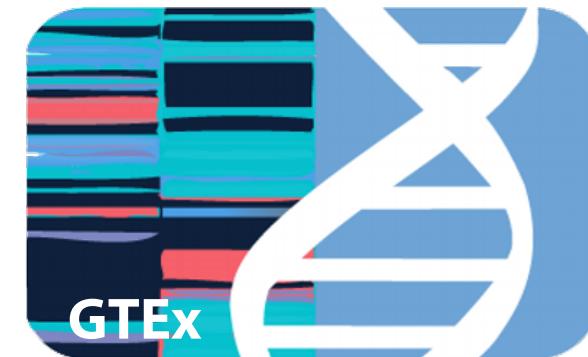
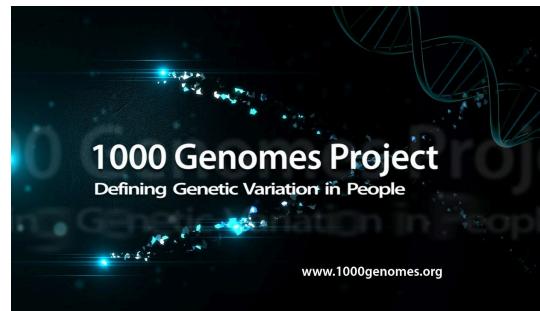
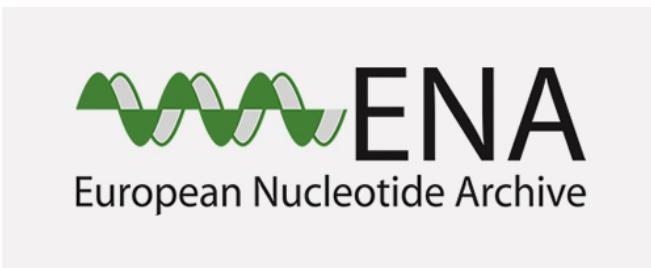
Google Index Estimate: >60 billion webpages

Google Universe Estimate (2013): >130 trillion webpages

Memory-Constrained Data Structures

What method would you use to build a search index on a collection of objects *in a memory-constrained environment?*

Constrained by Big Data (Large N)

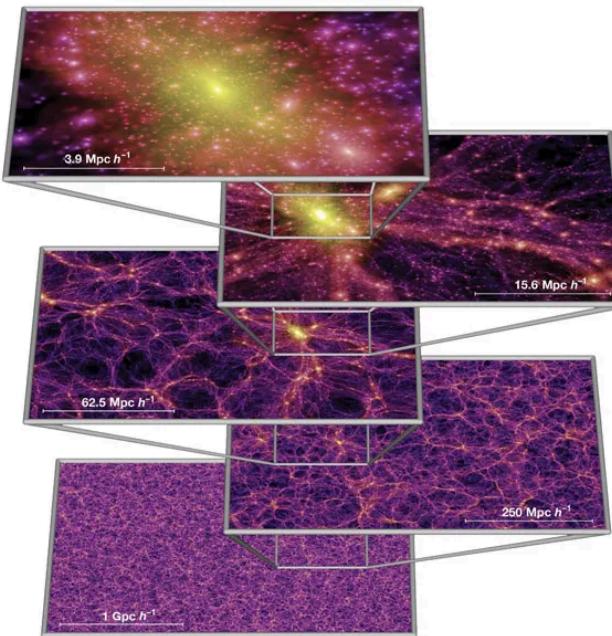
The logo for the Sequence Read Archive (SRA) is a dark blue rectangle. On the left, there is a small image of a sequencing gel showing DNA sequence data. To the right of the image, the acronym "SRA" is written in a large, white, sans-serif font. Below the acronym, a block of text provides a description of the SRA's purpose and the types of sequencing platforms it supports.

Sequence Read Archive Size: >60 petabases (10^{15})

Memory-Constrained Data Structures

What method would you use to build a search index on a collection of objects *in a memory-constrained environment?*

Constrained by Big Data (Large N)



Sky Survey Projects	Data Volume
DPOSS (The Palomar Digital Sky Survey)	3 TB
2MASS (The Two Micron All-Sky Survey)	10 TB
GBT (Green Bank Telescope)	20 PB
GALEX (The Galaxy Evolution Explorer)	30 TB
SDSS (The Sloan Digital Sky Survey)	40 TB
SkyMapper Southern Sky Survey	500 TB
PanSTARRS (The Panoramic Survey Telescope and Rapid Response System)	~ 40 PB expected
LSST (The Large Synoptic Survey Telescope)	~ 200 PB expected
SKA (The Square Kilometer Array)	~ 4.6 EB expected

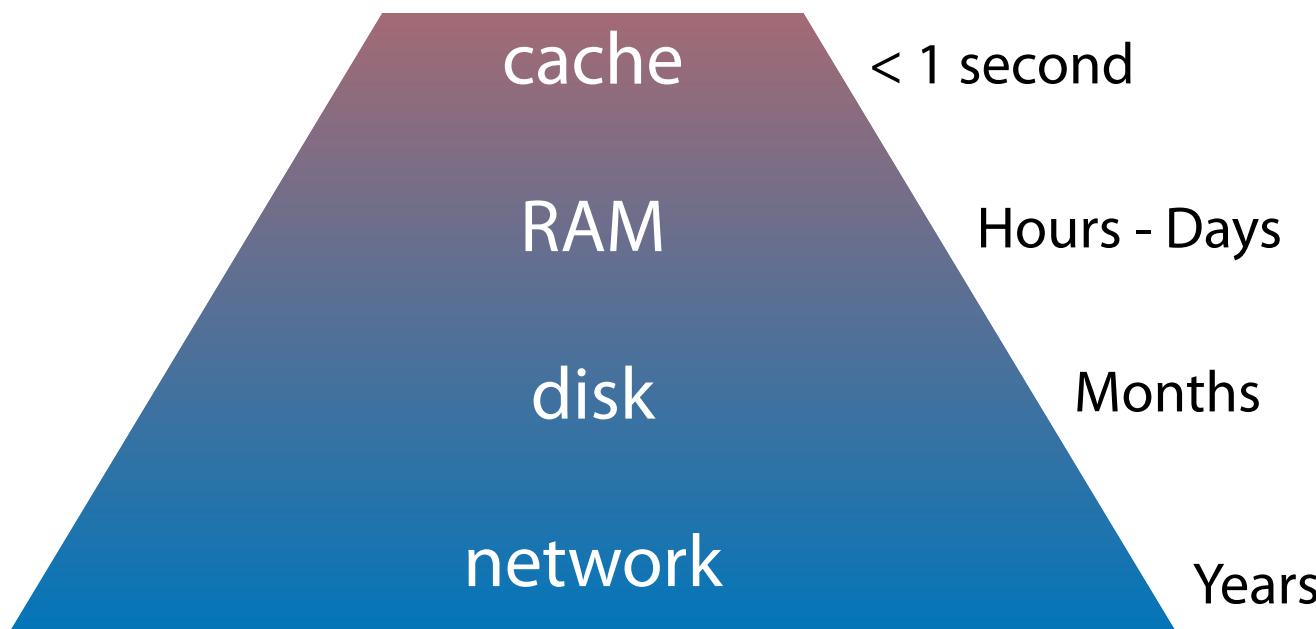
Table: <http://doi.org/10.5334/dsj-2015-011>

Estimated total volume of one array: 4.6 EB

Memory-Constrained Data Structures

What method would you use to build a search index on a collection of objects *in a memory-constrained environment?*

Constrained by resource limitations



(Estimates are Time x 1 billion courtesy of <https://gist.github.com/hellerbarde/2843375>)

Memory-Constrained Data Structures



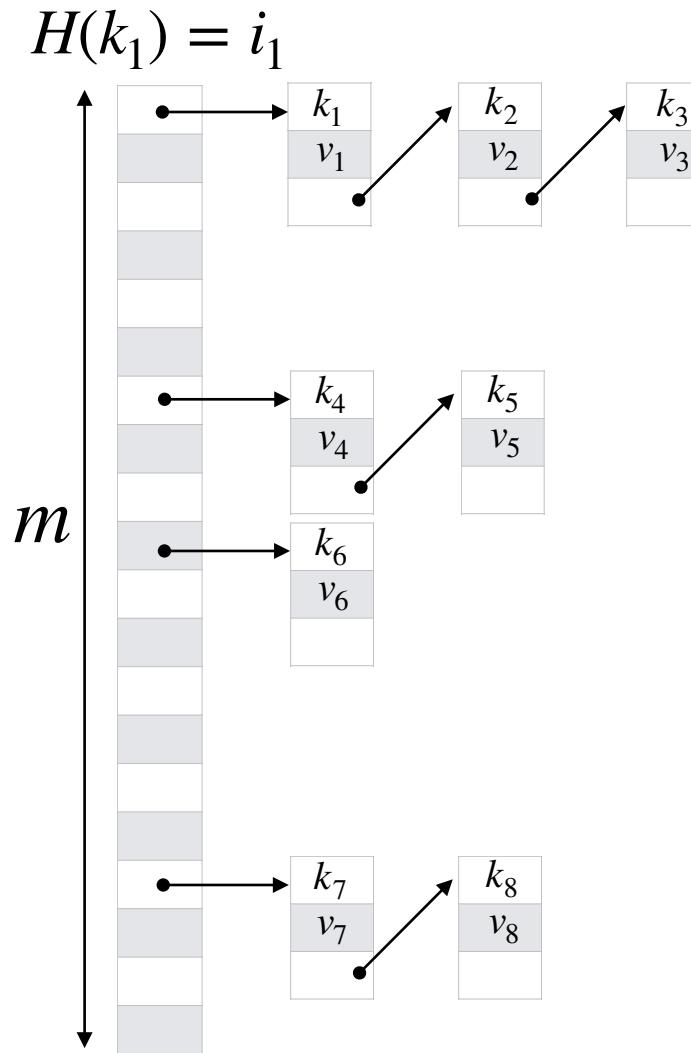
What method would you use to build a search index on a collection of objects *in a memory-constrained environment?*

Reducing storage costs

- 1) Throw out information that isn't needed
- 2) Compress the dataset

Reducing a hash table

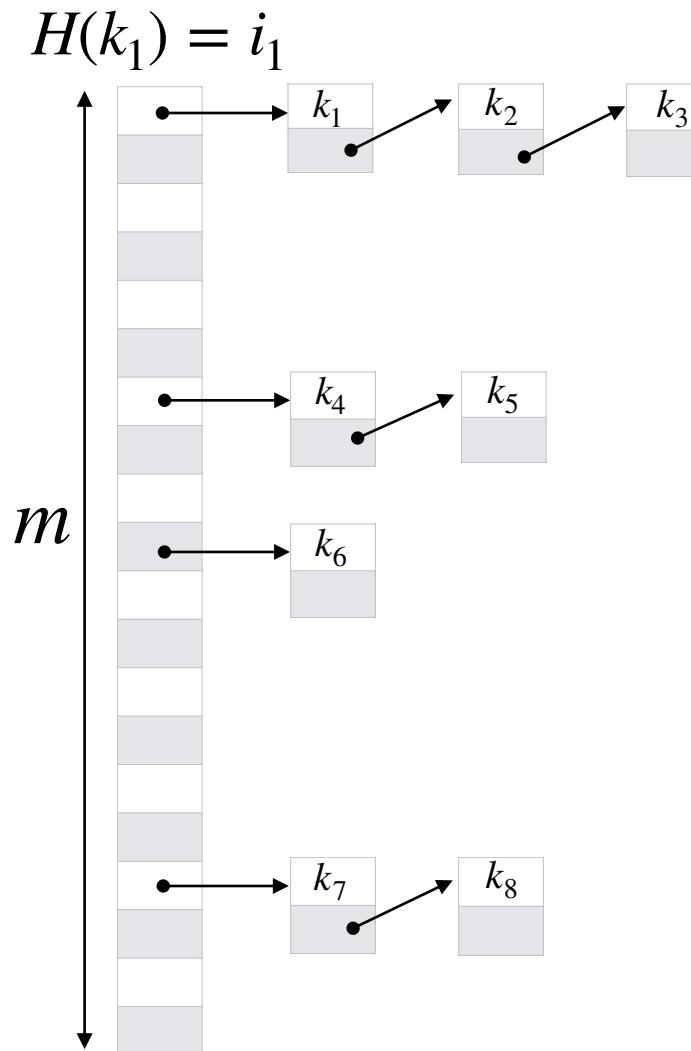
What can we remove from a hash table?



Reducing a hash table

What can we remove from a hash table?

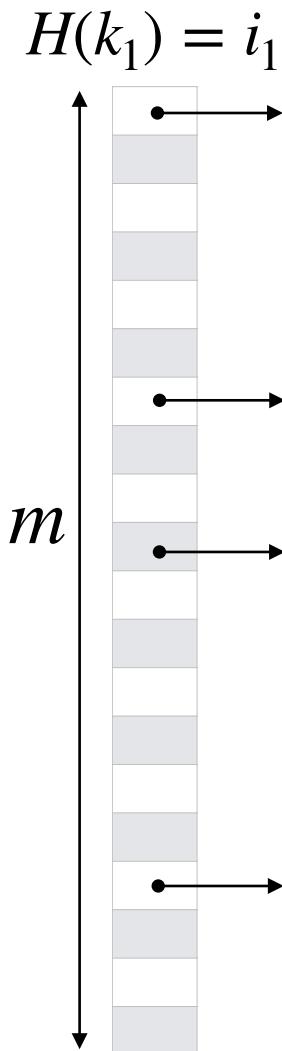
Take away values



Reducing a hash table

What can we remove from a hash table?

Take away values and keys



Reducing a hash table



What can we remove from a hash table?

Take away values and keys

This is a **bloom filter**

$$H(k_1) = i_1$$

1
0
0
0
0
0
1
0
0
1
0
0
0
0
0
0
1
0
0
0

Bloom Filter ADT

Constructor

Insert

Find

Bloom Filter: Insertion

$S = \{ 16, 8, 4, 13, 29, 11, 22 \}$

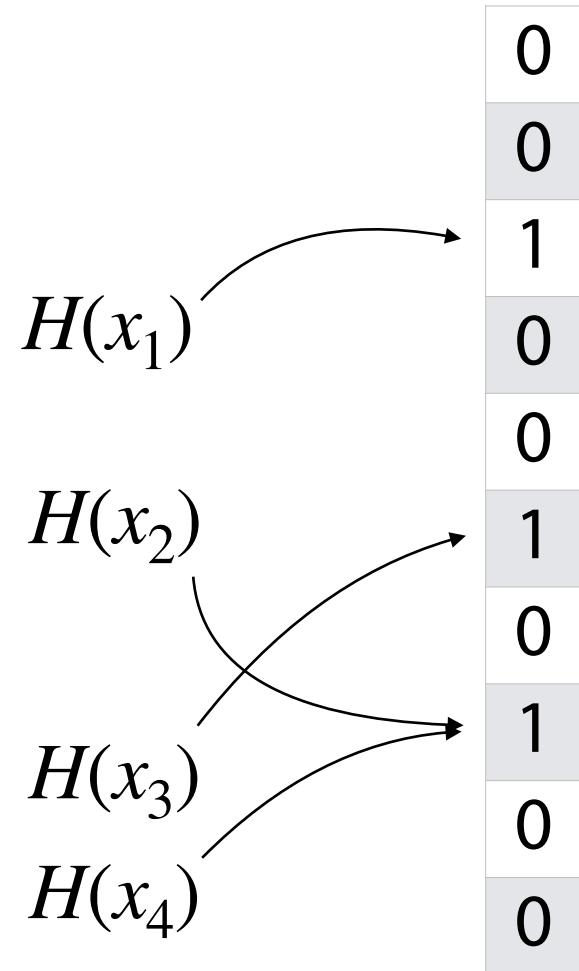
$h(k) = k \% 7$

0	0
1	0
2	0
3	0
4	0
5	0
6	0

Bloom Filter: Insertion

An item is inserted into a bloom filter by hashing and then setting the hash-valued bit to 1

If the bit was already one, it stays 1



Bloom Filter: Deletion

$S = \{ 16, 8, 4, 13, 29, 11, 22 \}$

_delete(13)

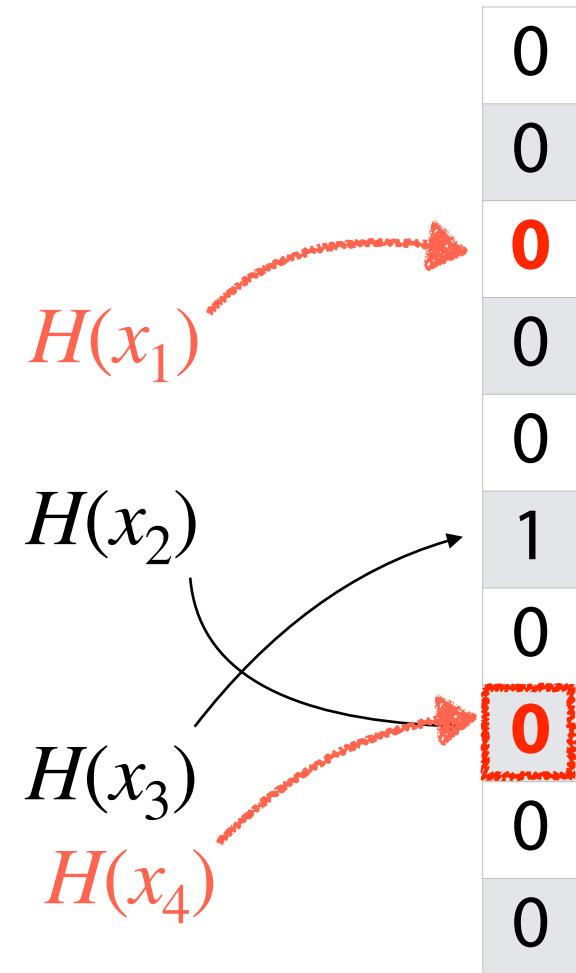
$h(k) = k \% 7$

0	0
1	1
2	1
3	0
4	1
5	0
6	1

_delete(29)

Bloom Filter: Deletion

Due to hash collisions and lack of information, items cannot be deleted!



Bloom Filter: Search

$S = \{ 16, 8, 4, 13, 29, 11, 22 \}$

$h(k) = k \% 7$

0	0
1	1
2	1
3	0
4	1
5	0
6	1

$_find(16)$

$_find(20)$

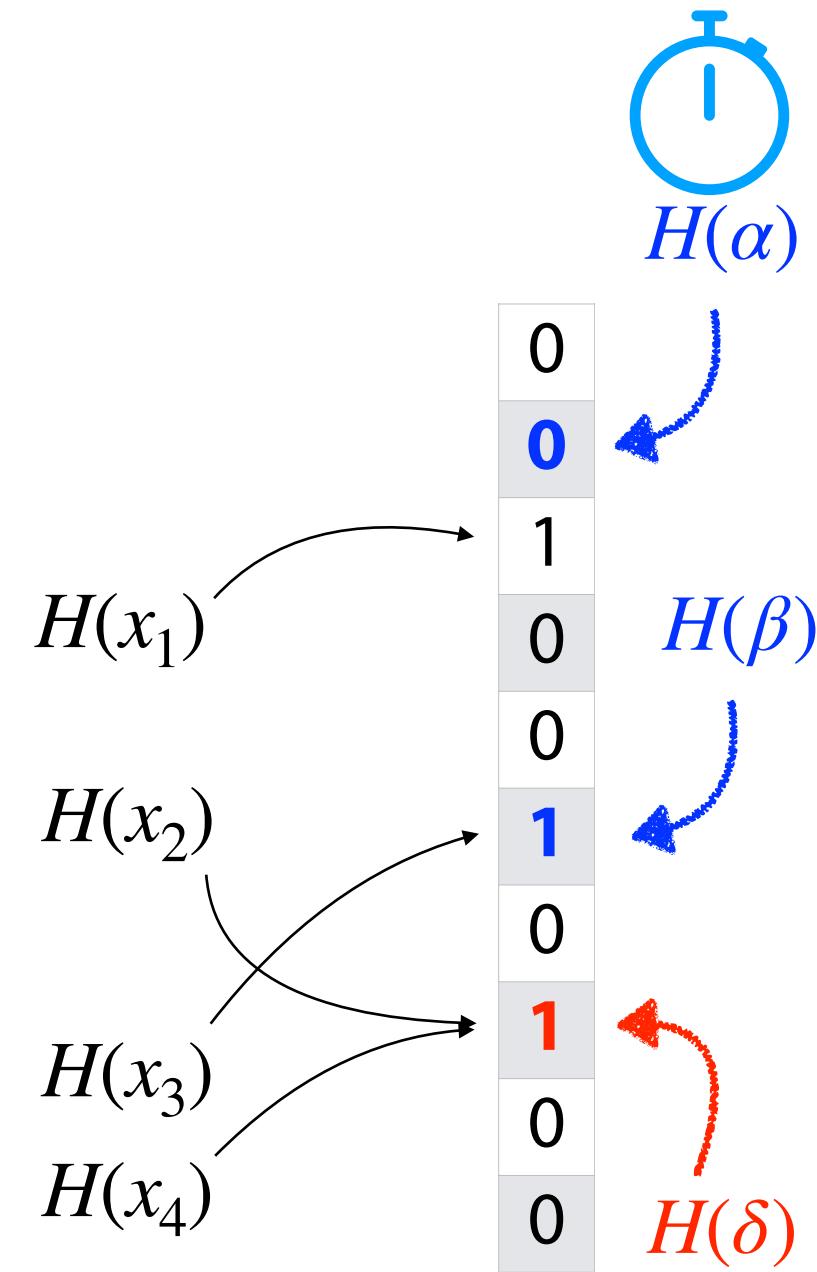
$_find(3)$

Bloom Filter: Search

The bloom filter is a *probabilistic* data structure!

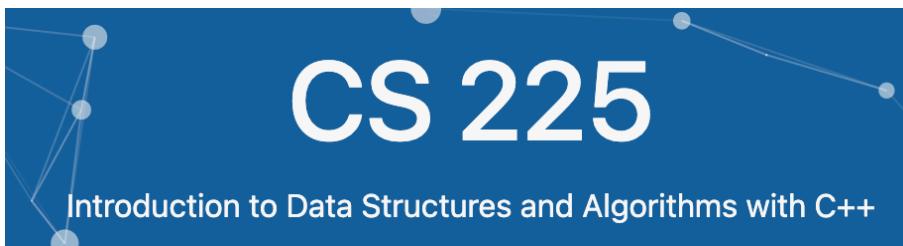
If the value in the BF is 0:

If the value in the BF is 1:

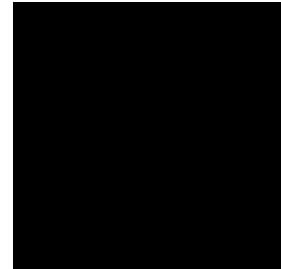
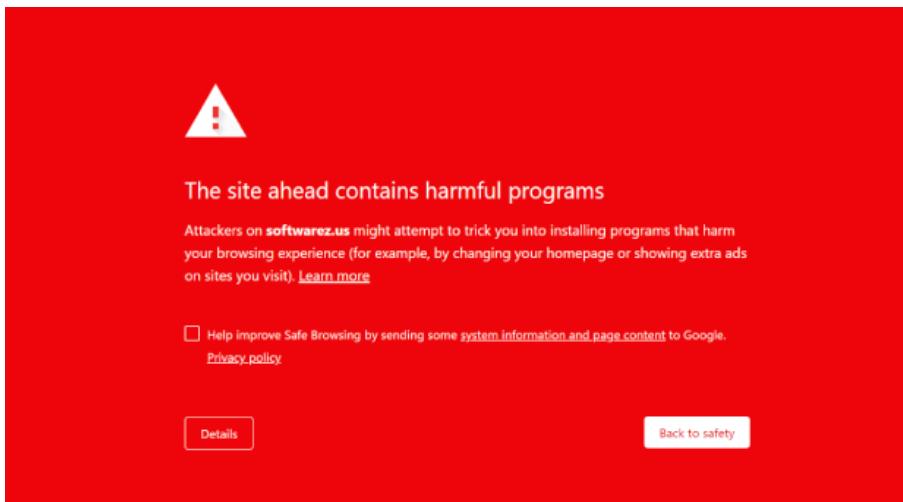


Probabilistic Accuracy: Malicious Websites

Imagine we have a detection oracle that identifies if a site is malicious



“Not malicious”



“Malicious”

Probabilistic Accuracy: Malicious Websites

Imagine we have a detection oracle that identifies if a site is malicious

True Positive:

False Positive:

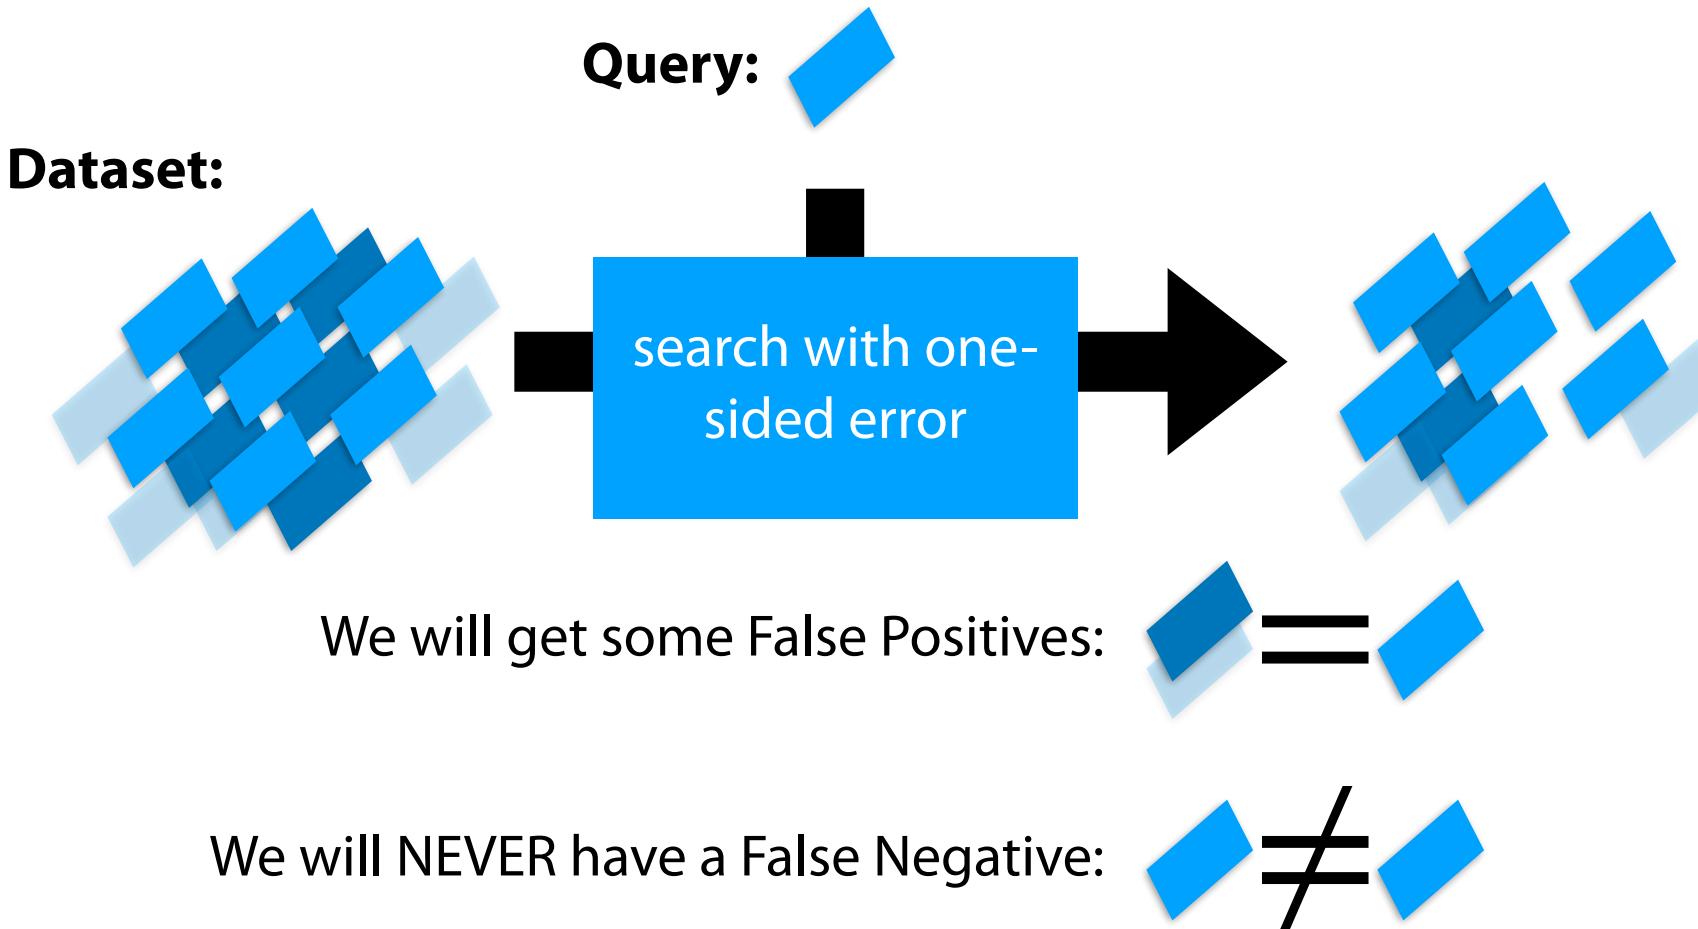
False Negative:

True Negative:

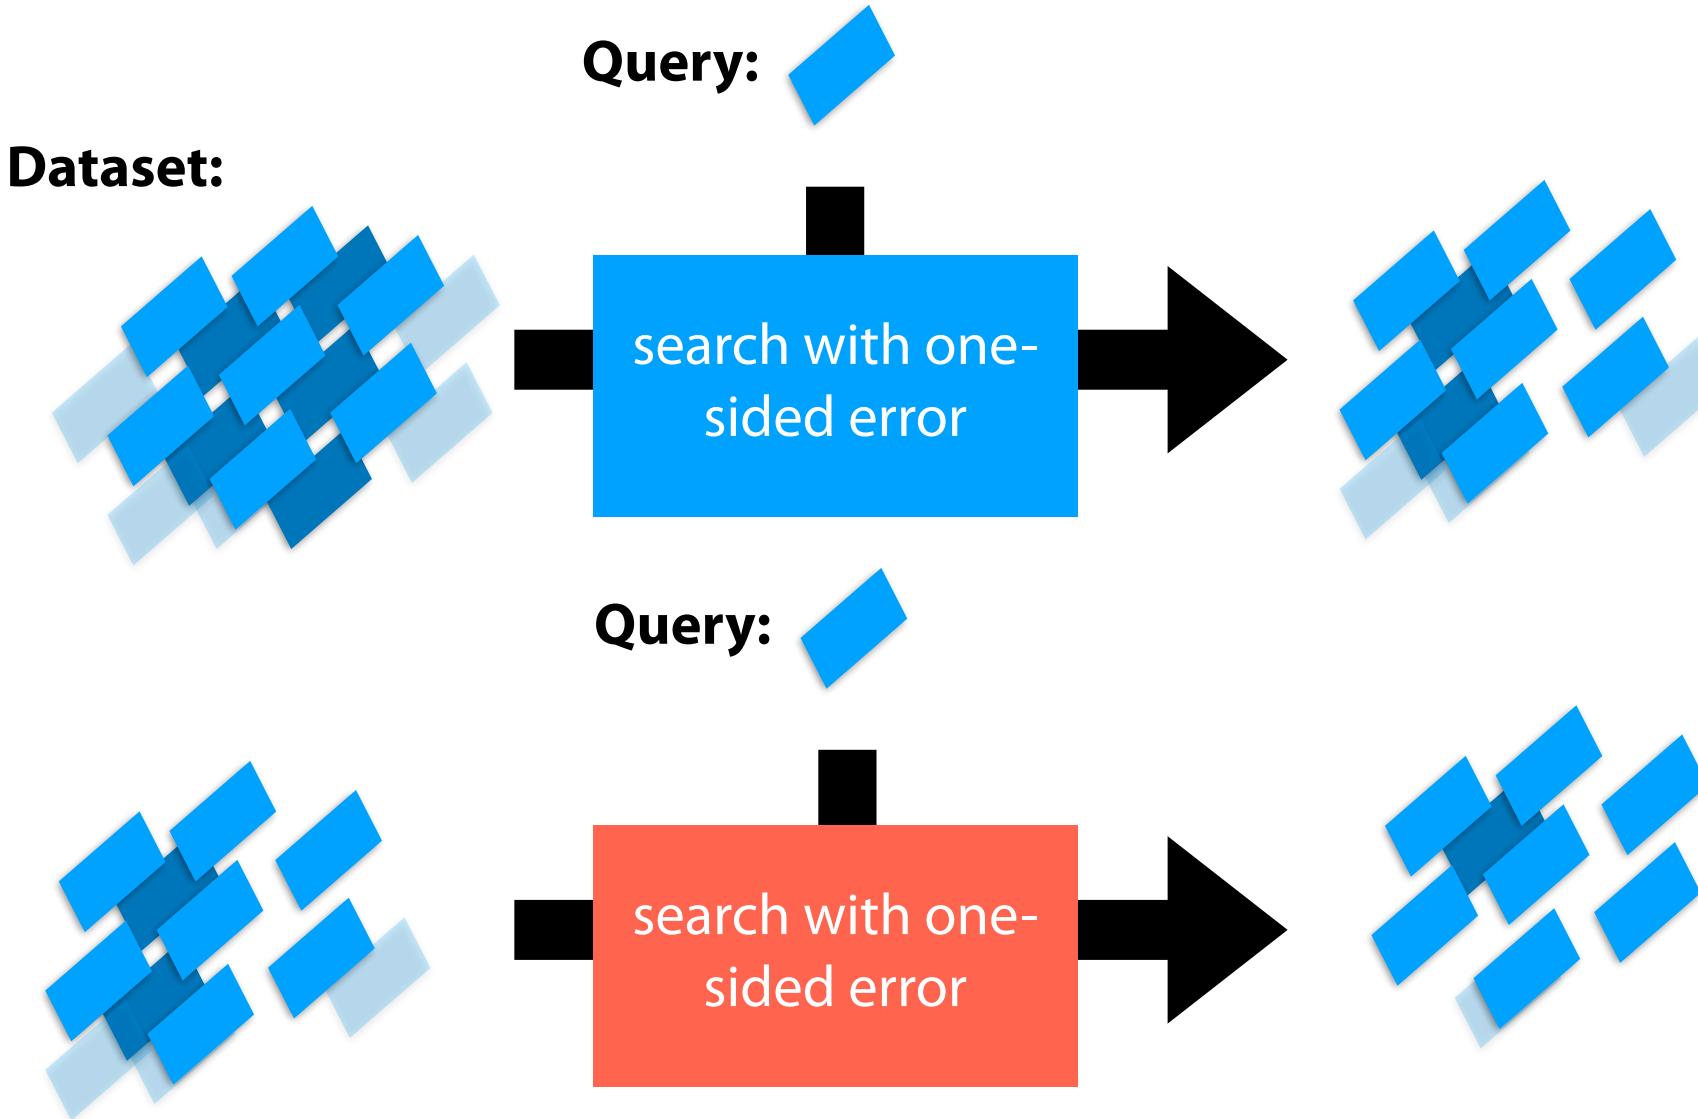
Imagine we have a **bloom filter** that **stores malicious sites**...

		Bit Value = 1	Bit Value = 0									
		$H(z)$	$H(z)$									
Item Inserted	True Positive	<table border="1"><tr><td>0</td></tr><tr><td>1</td></tr><tr><td>0</td></tr><tr><td>0</td></tr><tr><td>1</td></tr></table> 'Yes'	0	1	0	0	1	False Negative				
0												
1												
0												
0												
1												
False Positive	<table border="1"><tr><td>0</td></tr><tr><td>1</td></tr><tr><td>0</td></tr><tr><td>0</td></tr><tr><td>1</td></tr></table> 'Yes'	0	1	0	0	1	<table border="1"><tr><td>0</td></tr><tr><td>0</td></tr><tr><td>0</td></tr><tr><td>0</td></tr><tr><td>1</td></tr></table> 'No'	0	0	0	0	1
0												
1												
0												
0												
1												
0												
0												
0												
0												
1												
Item NOT inserted	False Positive	<table border="1"><tr><td>0</td></tr><tr><td>1</td></tr><tr><td>0</td></tr><tr><td>0</td></tr><tr><td>1</td></tr></table> 'Yes'	0	1	0	0	1	True Negative				
0												
1												
0												
0												
1												
True Negative	<table border="1"><tr><td>0</td></tr><tr><td>0</td></tr><tr><td>0</td></tr><tr><td>0</td></tr><tr><td>1</td></tr></table> 'No'	0	0	0	0	1						
0												
0												
0												
0												
1												

Probabilistic Accuracy: One-sided error



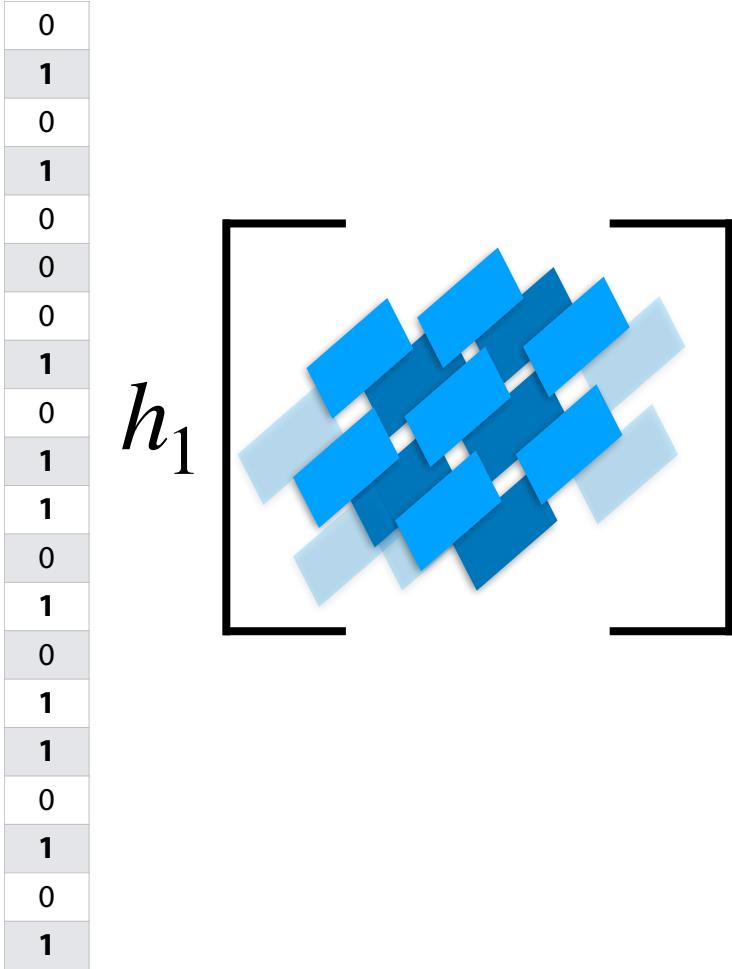
Probabilistic Accuracy: One-sided error



...

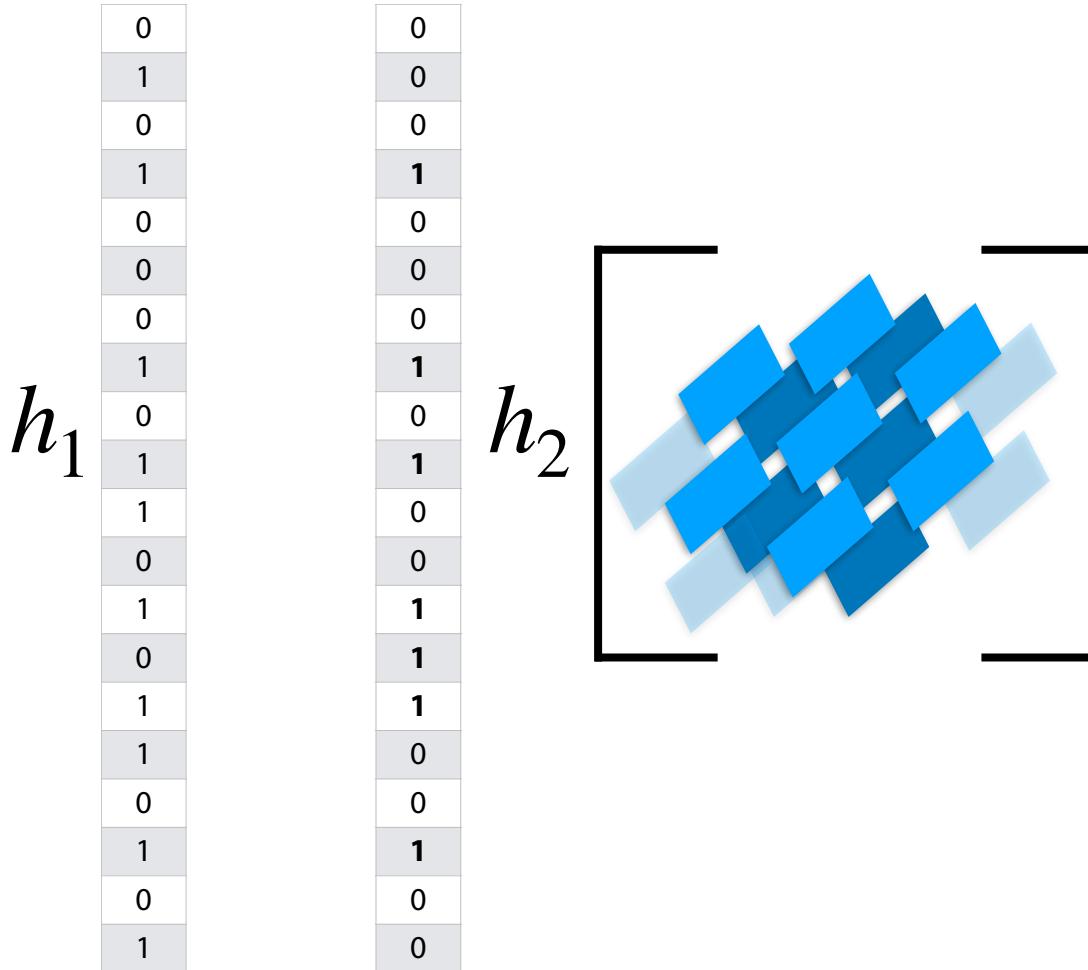
Bloom Filter: Repeated Trials

Use many hashes/filters; add each item to each filter



Bloom Filter: Repeated Trials

Use many hashes/filters; add each item to each filter



Bloom Filter: Repeated Trials

Use many hashes/filters; add each item to each filter

0
1
0
1
0
0
0
1
0
1
1
0
1
0
1
1
0
1
0
1

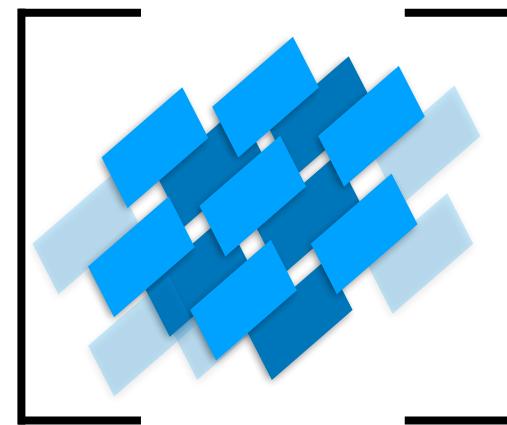
h_1

0
0
0
1
0
0
0
1
0
1
0
1
0
1
1
0
0
1
0

h_2

0
1
1
0
0
0
1
0
1
0
1
1
0
1
1
0
0
1
0

h_3



Bloom Filter: Repeated Trials

Use many hashes/filters; add each item to each filter

0	0	0	0
1	0	1	1
0	0	1	1
1	1	1	1
0	0	0	1
0	0	0	1
0	0	1	0
1	1	1	0
0	0	0	0
1	1	1	0
1	0	1	0
0	0	0	1
1	1	1	0
0	1	0	0
1	0	1	1
0	0	0	1
1	1	1	1
0	0	0	1
1	1	1	1
0	0	0	1
1	1	1	1

Bloom Filter: Repeated Trials

0
1
0
1
0
0
0
1
0
1
1
0
1
0
1
1
0
1
0
1
0
1

0
0
0
1
0
0
0
1
1
0
0
1
1
0
0
1
0
1
0
0
0
0

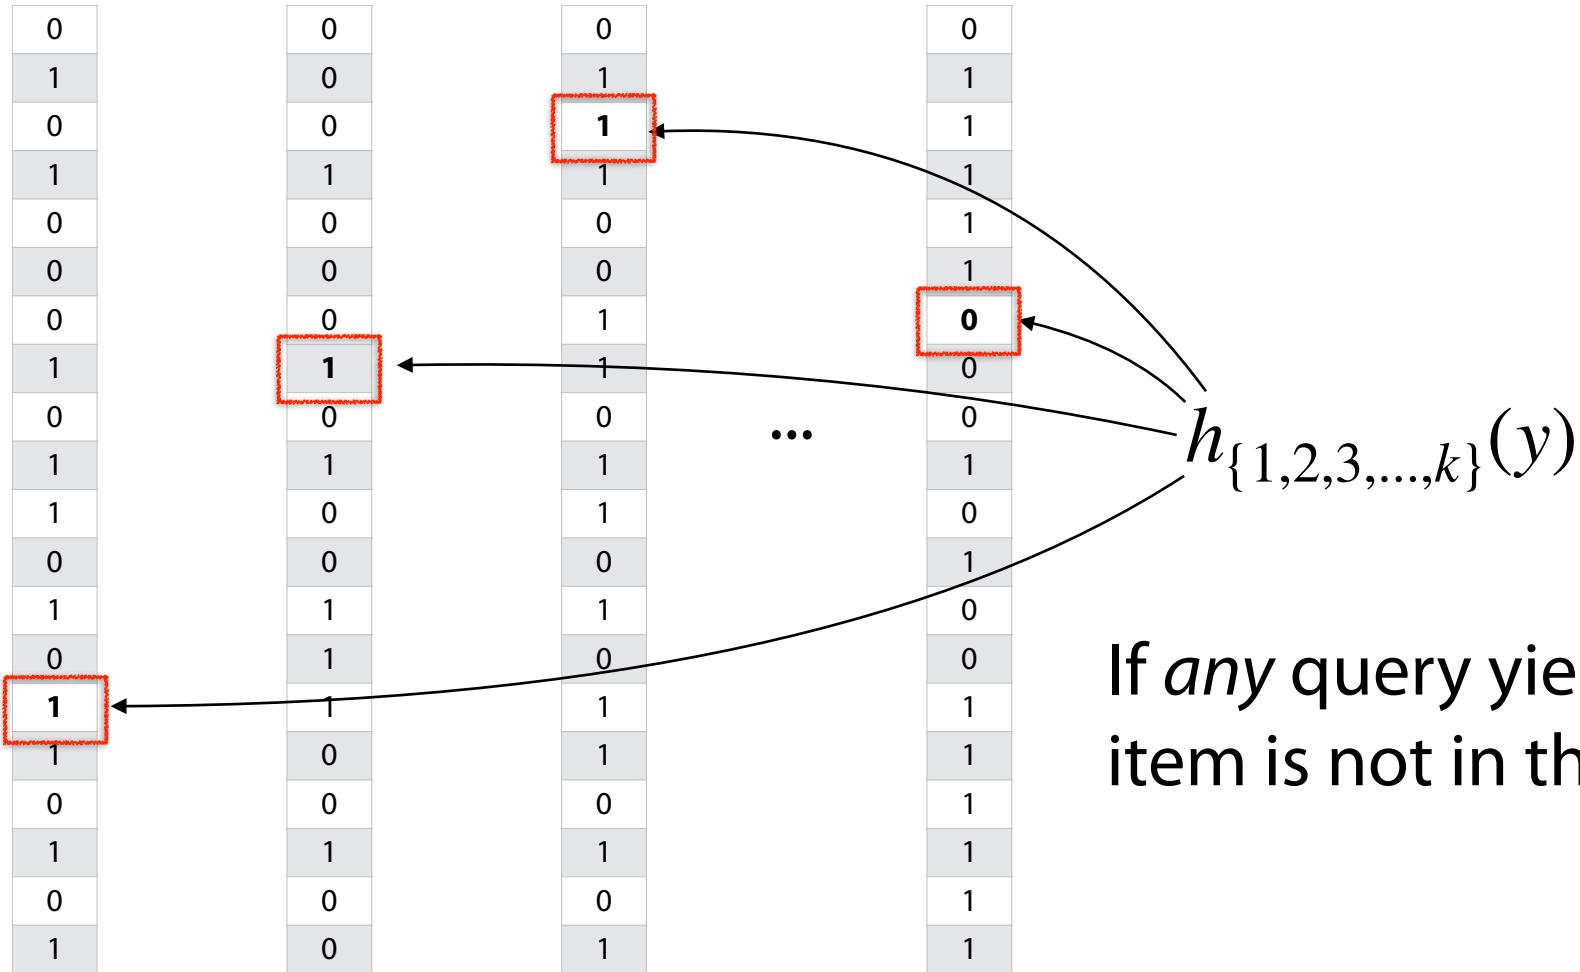
0
1
1
1
0
0
0
1
1
0
0
1
1
0
1
1
0
0
1
0
1

...

0
1
1
1
1
0
0
0
0
0
1
1
1
1
1
1
1
1
1
1
1

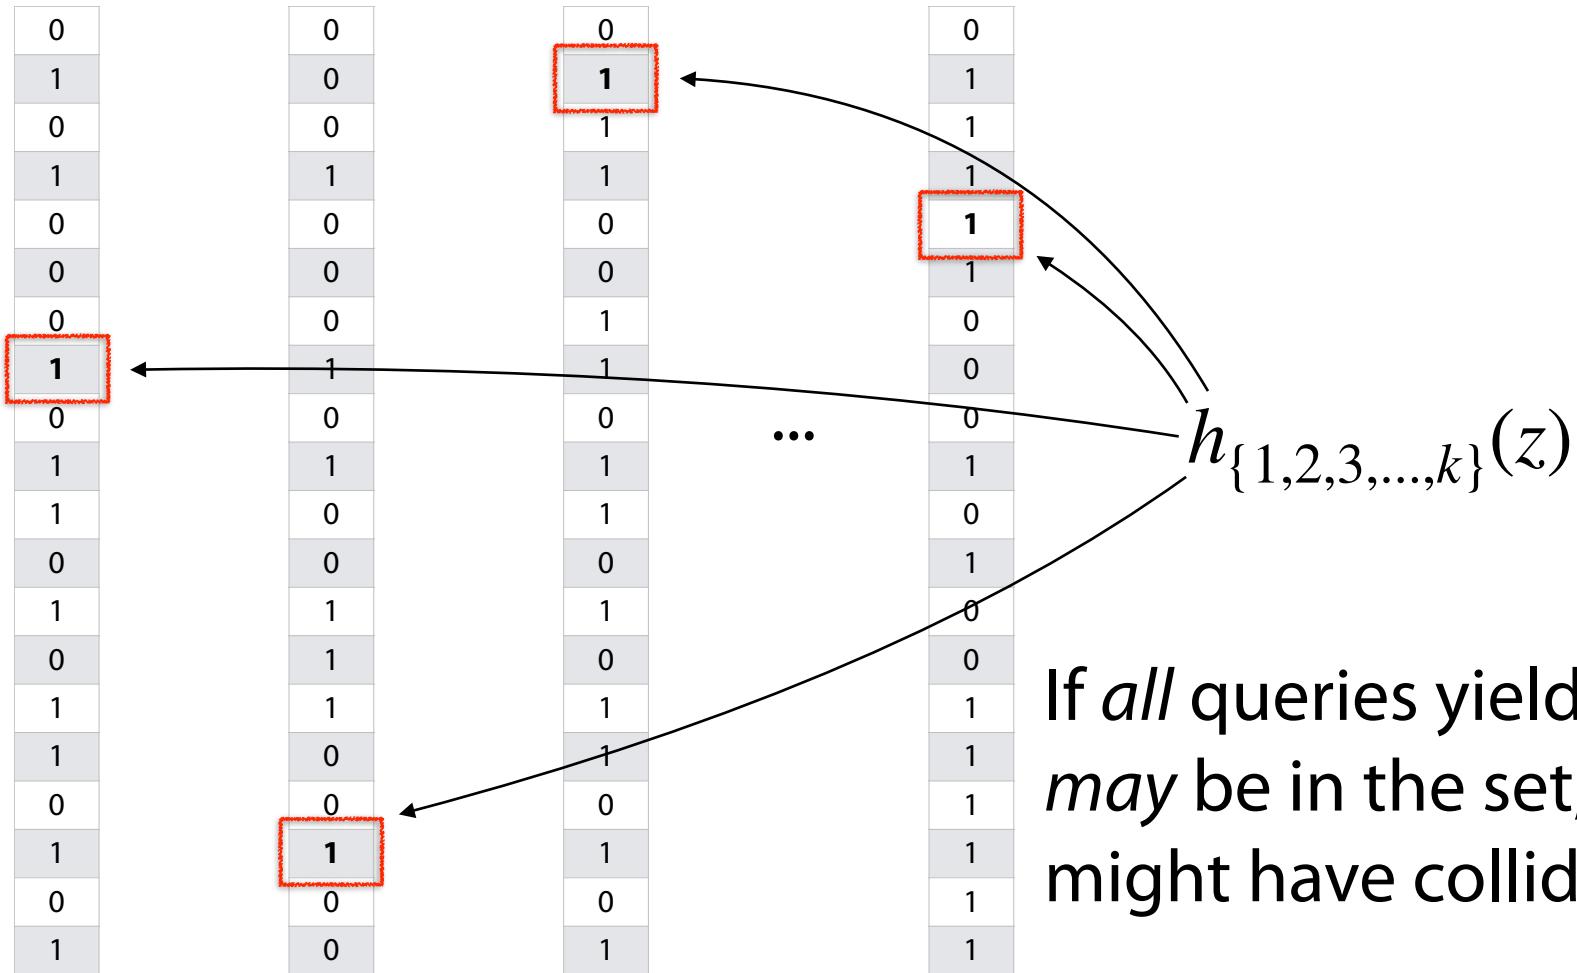
$h_{\{1,2,3,\dots,k\}}(y)$

Bloom Filter: Repeated Trials



If *any* query yields 0,
item is not in the set

Bloom Filter: Repeated Trials



If *all* queries yield 1, item *may* be in the set; or we might have collided *k* times

Bloom Filter: Repeated Trials

Using repeated trials, even a very bad filter can still have a very low FPR!

If we have k bloom filter, each with a FPR p , what is the likelihood that ***all*** filters return the value '1' for an item we didn't insert?

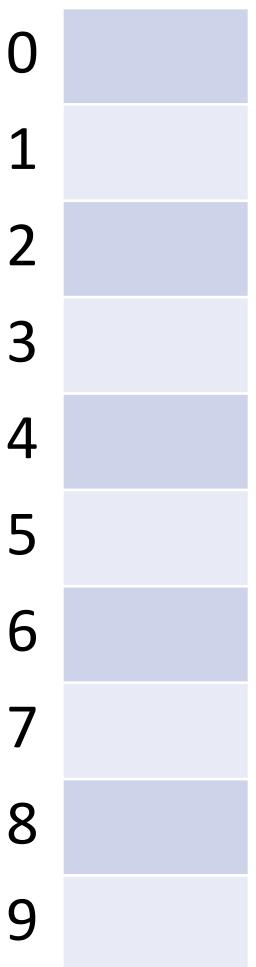
Bloom Filter: Repeated Trials

But doesn't this hurt our storage costs by storing k separate filters?

h_1	h_2	h_3	...	h_k
0	0	0		0
1	0	1		1
0	0	1		1
1	1	1		1
0	0	0		1
0	0	0		1
0	0	1		0
1	1	1		0
0	0	0		0
1	1	1		1
1	0	1		0
0	0	1		0
1	1	0		1
0	0	1		1
1	0	0		1
0	0	0		1
1	1	1		1
0	0	0		1
1	0	1		1
0	0	0		1
1	1	1		1

Bloom Filter: Repeated Trials

Rather than use a new filter for each hash, one filter can use k hashes



$$S = \{ 6, 8, 4 \}$$

1

$$h_1(x) = x \% 10$$

2

$$h_2(x) = 2x \% 10$$

3

$$h_3(x) = (5+3x) \% 10$$

4

5

6

7

8

9

Bloom Filter: Repeated Trials

Rather than use a new filter for each hash, one filter can use k hashes

0	0	$h_1(x) = x \% 10$	$h_2(x) = 2x \% 10$	$h_3(x) = (5+3x) \% 10$
1	0			
2	1	<u>find(1)</u>		
3	1			
4	1			
5	0			
6	1	<u>find(16)</u>		
7	1			
8	1			
9	1			

Bloom Filter



$$H = \{h_1, h_2, \dots, h_k\}$$

A probabilistic data structure storing a set of values

Built from a bit vector of length m and k hash functions

Insert / Find runs in: _____

Delete is not possible (yet)!

0
0
1
0
0
0
1
0
1
0
0
0