

### What is a Filter?



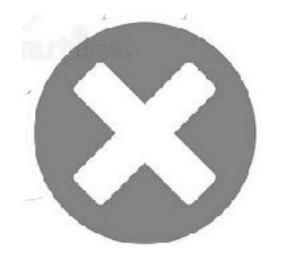
#### What is a Filter?

Does X exist? → X Y Z

#### What is a Filter?

Does X exist?  $\longrightarrow$  X Y Z

No false negatives



false positives with tunable probability



## Why use a Filter?

Data

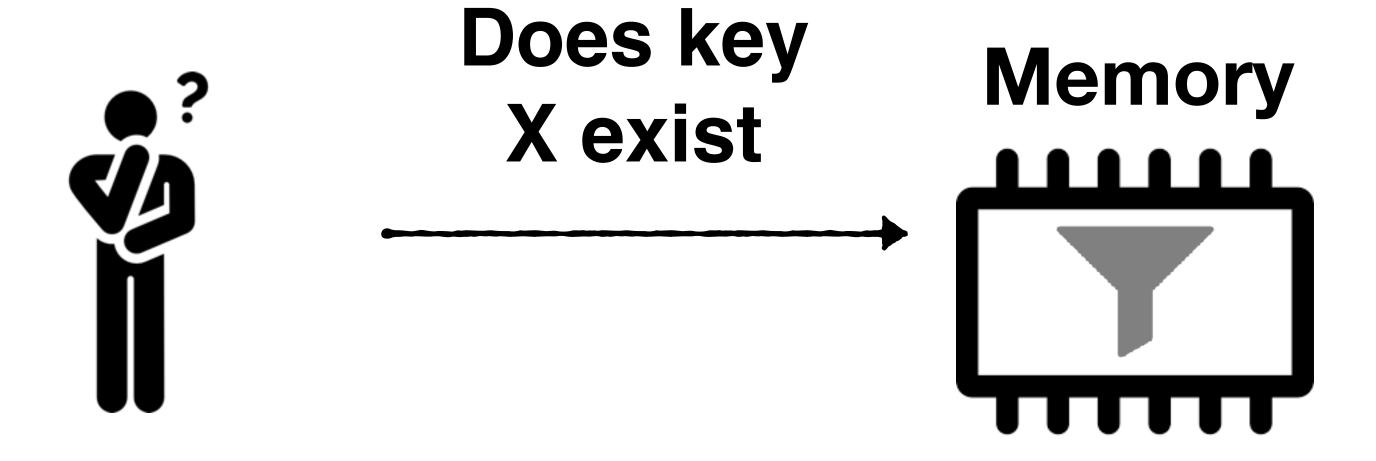




# Does key X exist

## Data

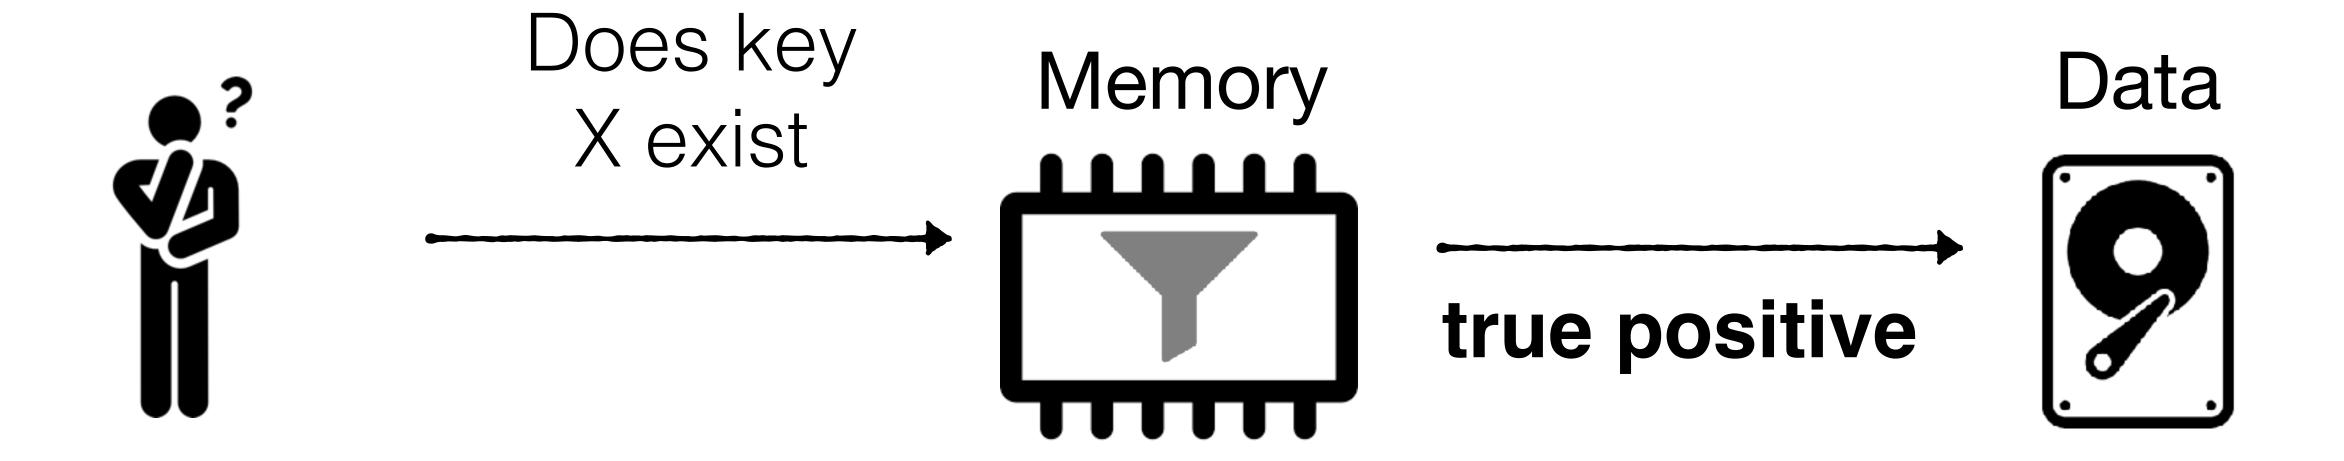


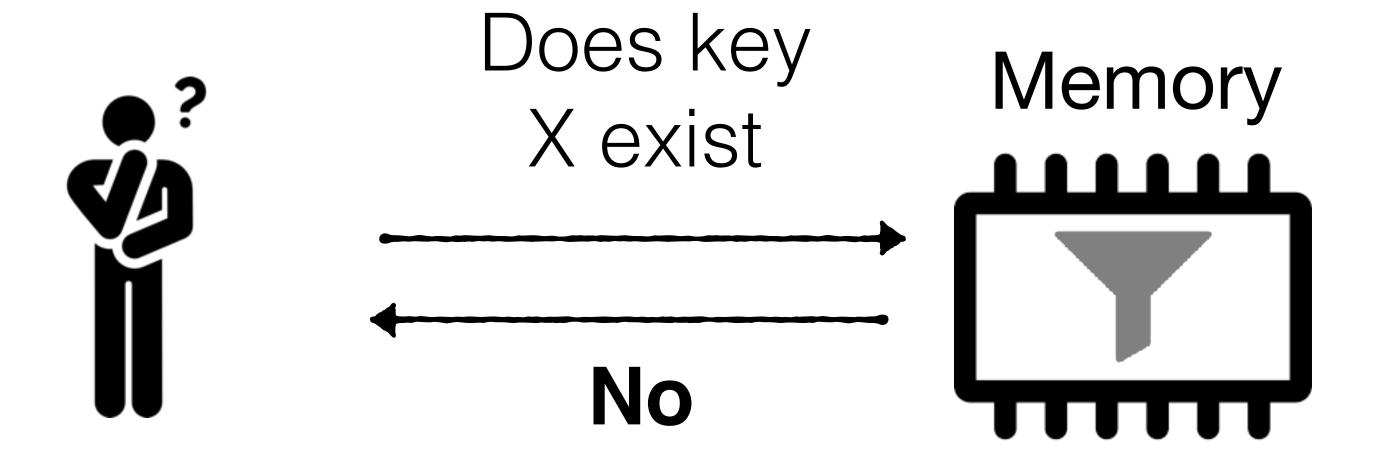


Data

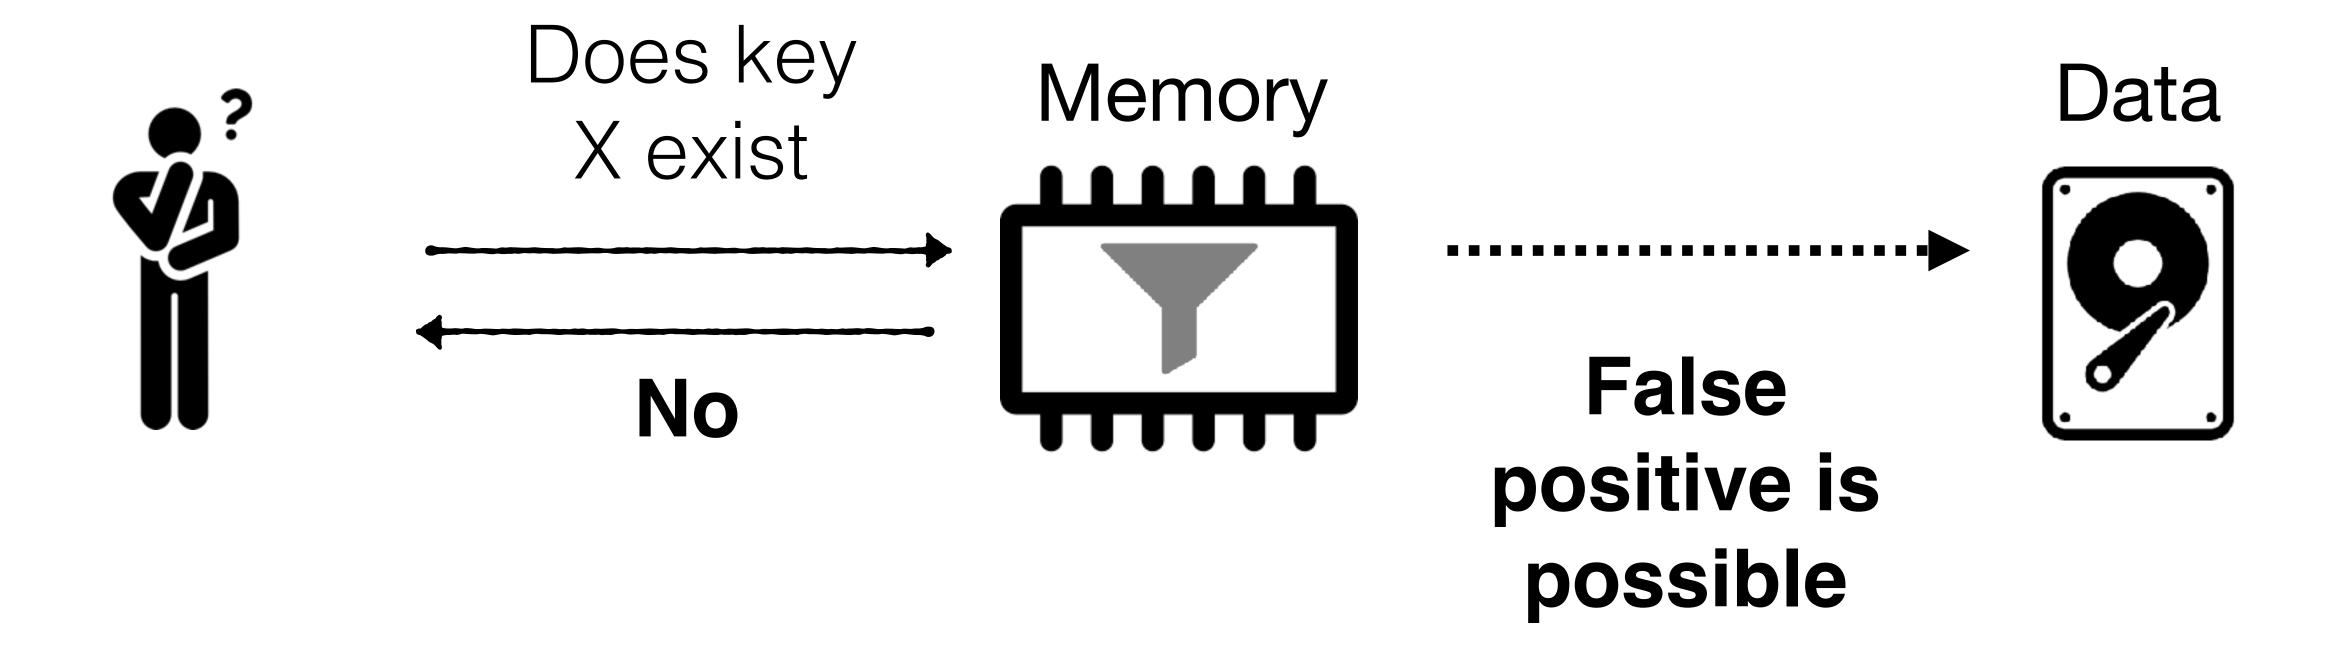


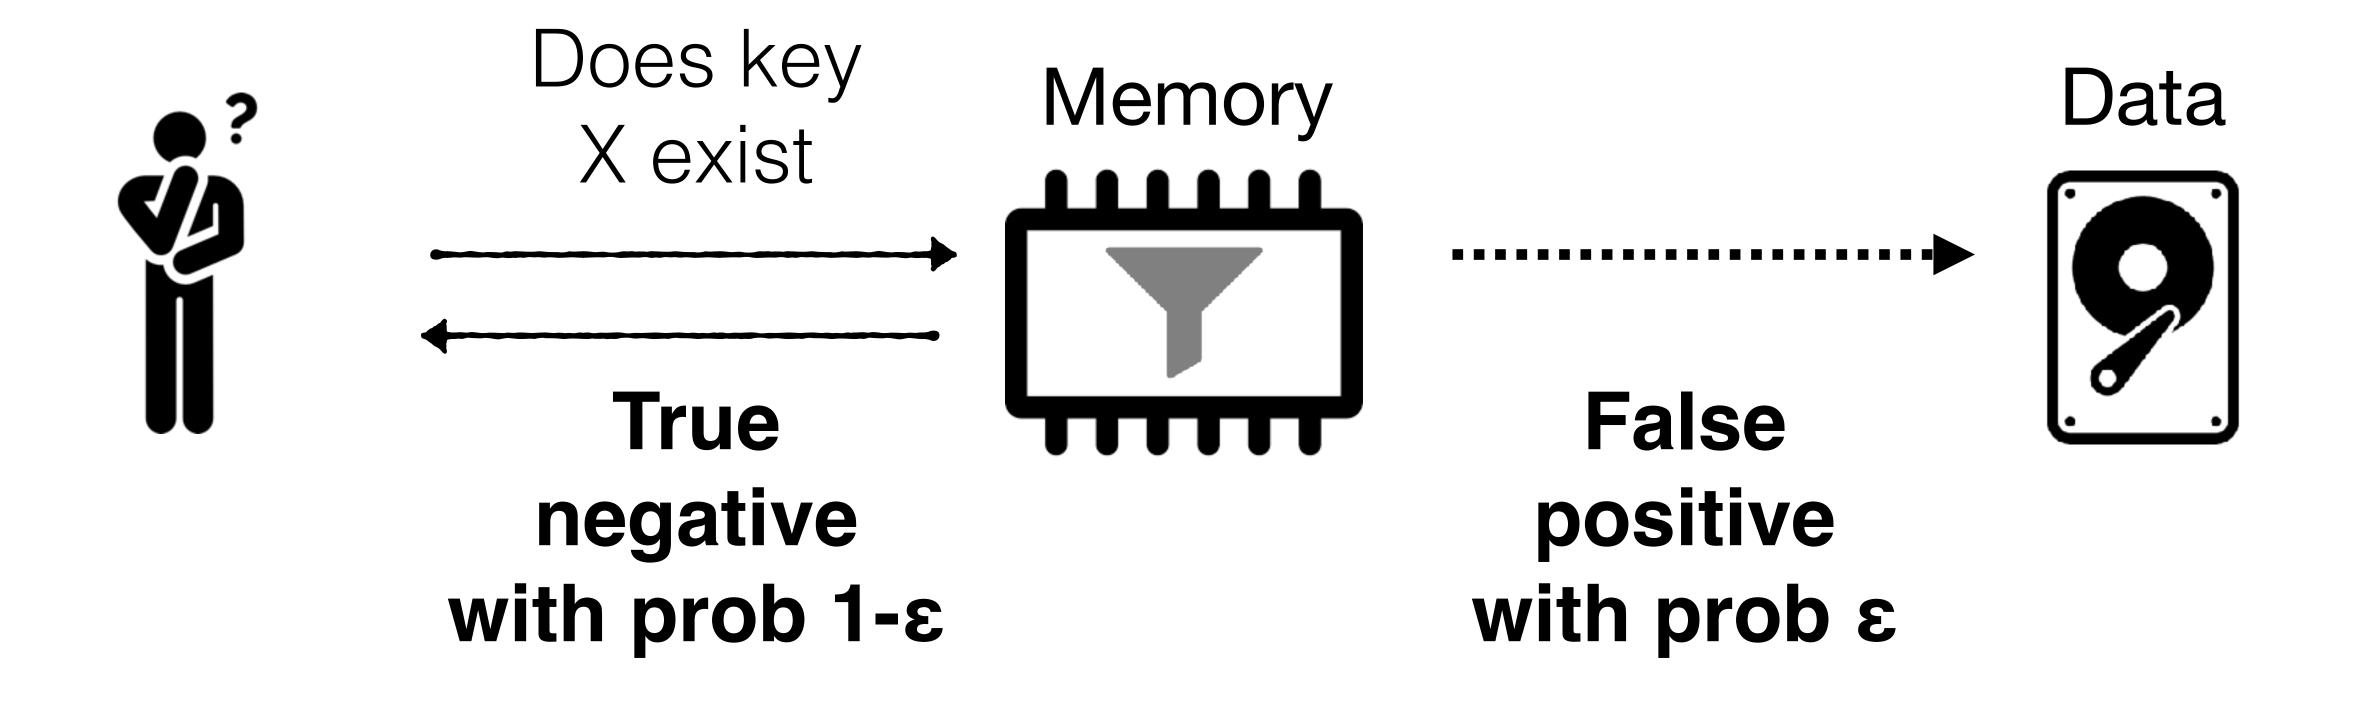
## If key X exists

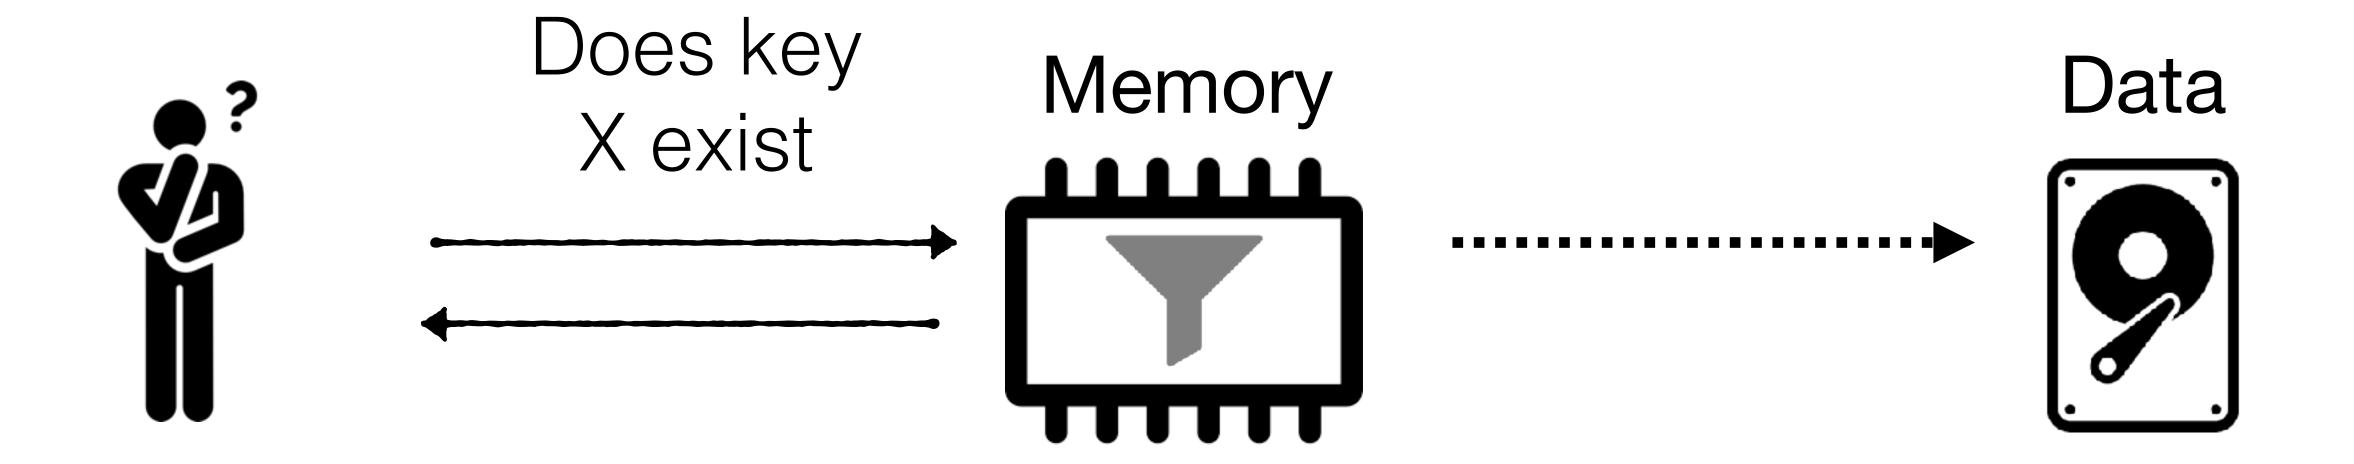












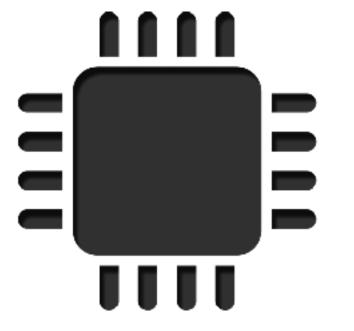
ε - false positive rate - FPR

## Why care about filters?

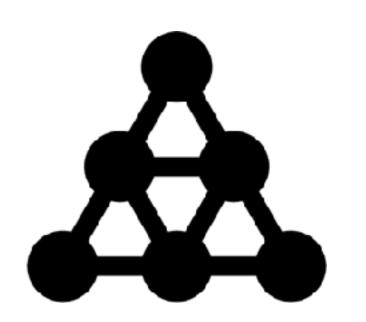
# Widely used in systems



Hardware Optimizations



Algorithmic Reasoning/ Techniques



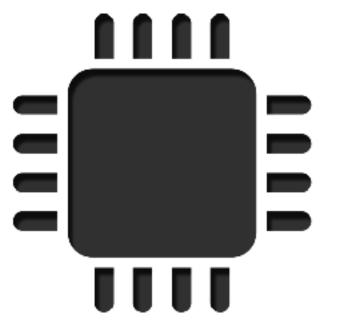
## Why care about filters?

#### Means to learn

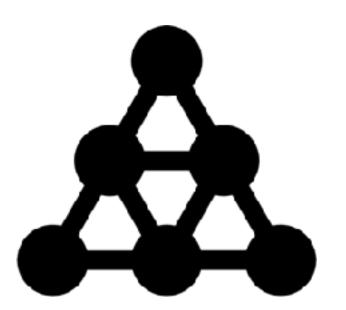
Widely used in systems



Hardware Optimizations

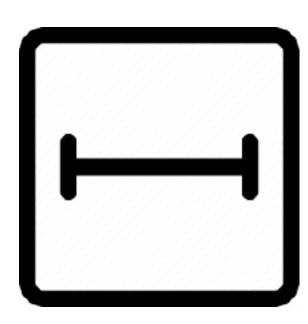


Algorithmic Reasoning/ Techniques





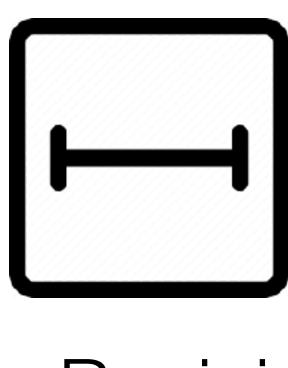
No deletes



No Resizing

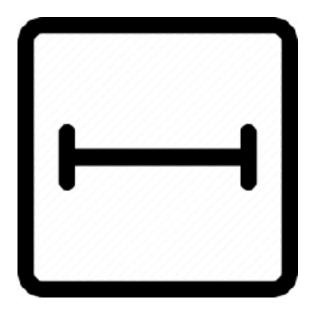


No deletes

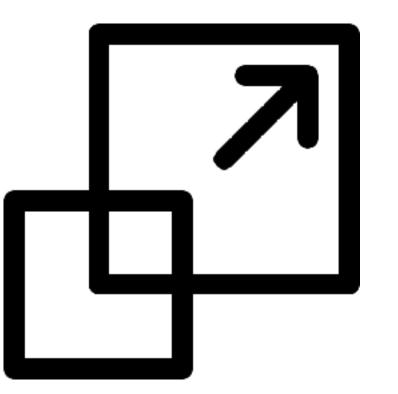


No Resizing

Modifications require rebuilding from scratch

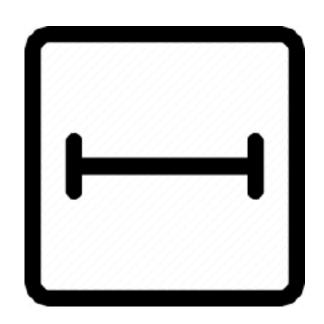


# Dynamic Filters (next week)



Delete + Resize

Dynamic Filters

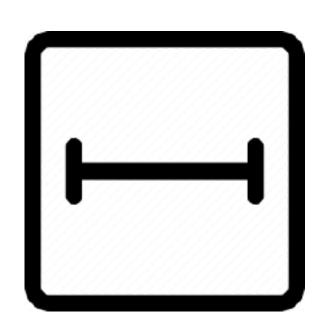


**Fastest Queries** 

Lowest FPR

Delete + Resize

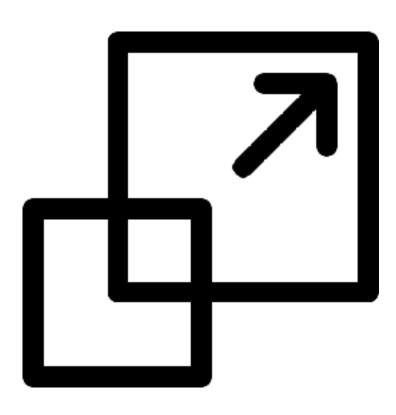
Dynamic Filters



Fastest Queries Lowest

FPR

But not both at same time



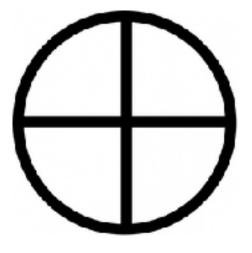
Delete + Resize

Bloom



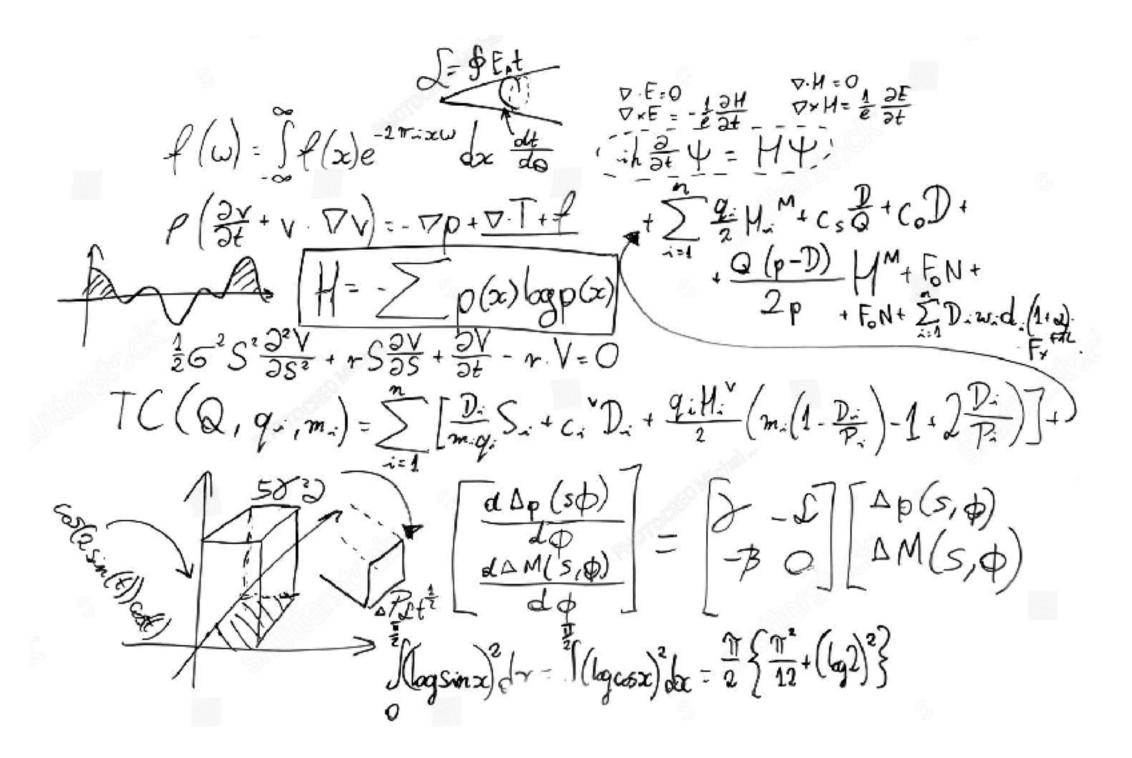
**Fastest Queries** 

XOR



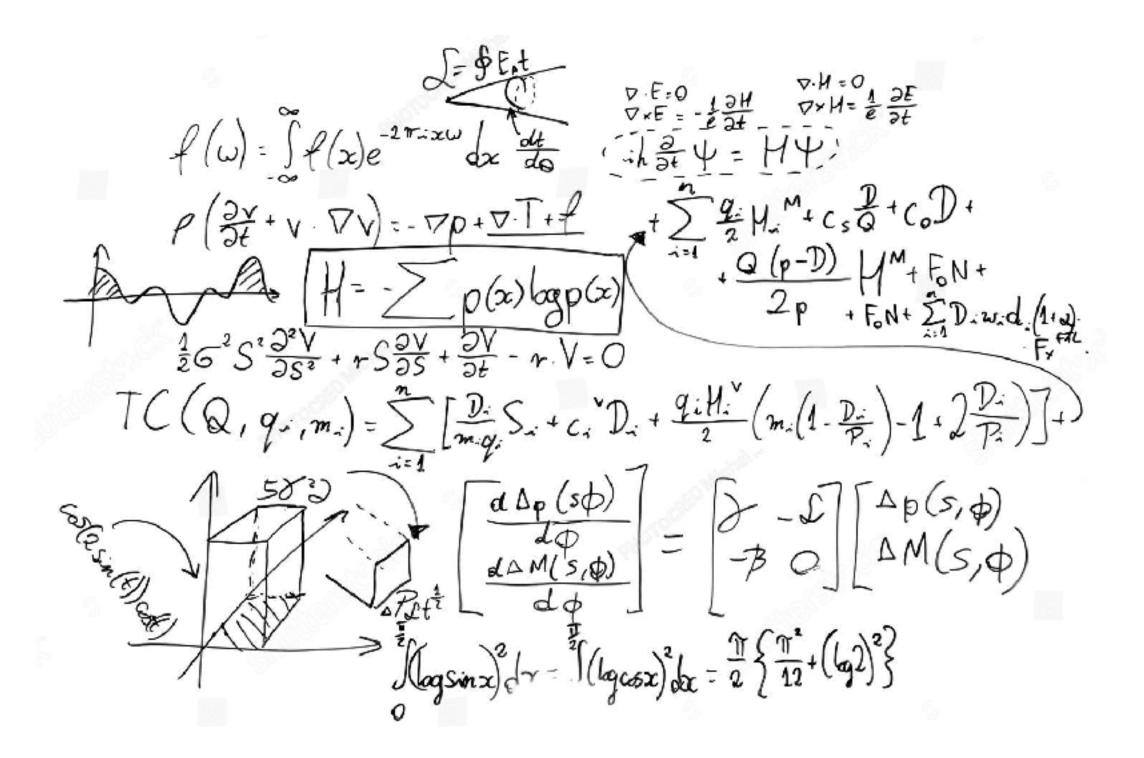
Lowest FPR

#### Note



Today is more mathematical than usual

#### Note



Today is more mathematical than usual

(Do not be intimidated)

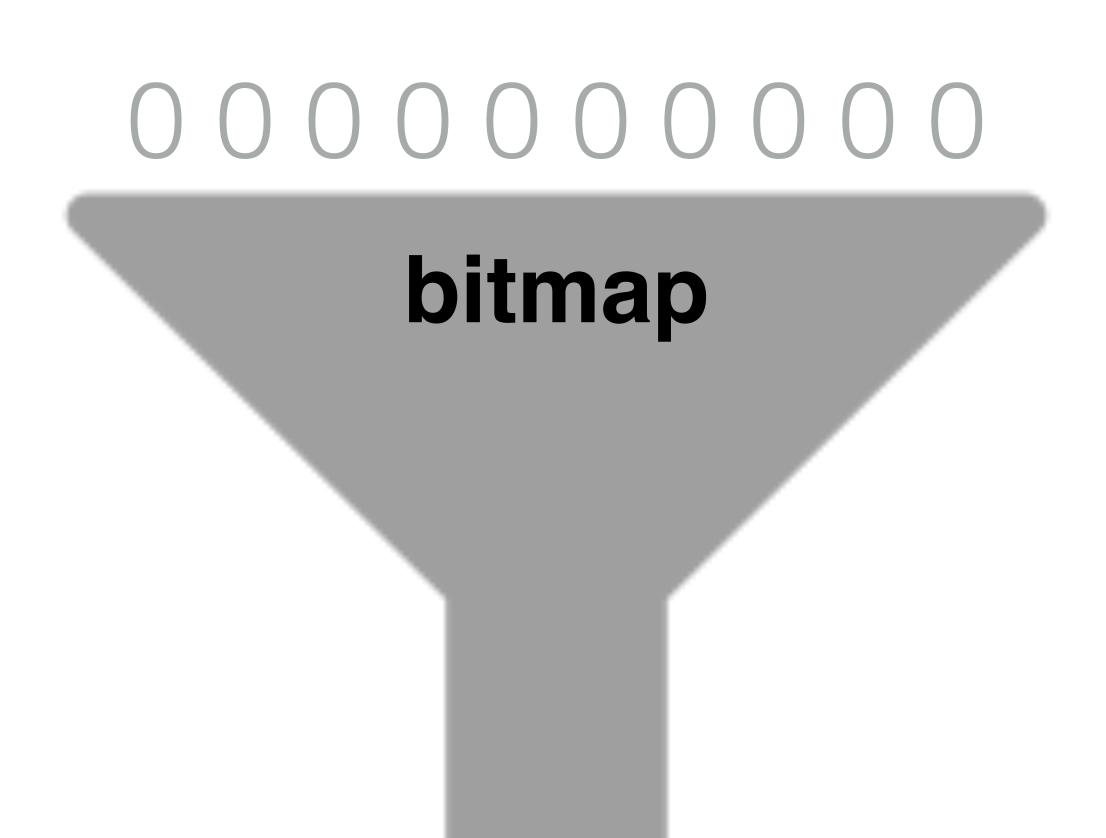
# **Bloom Filters**



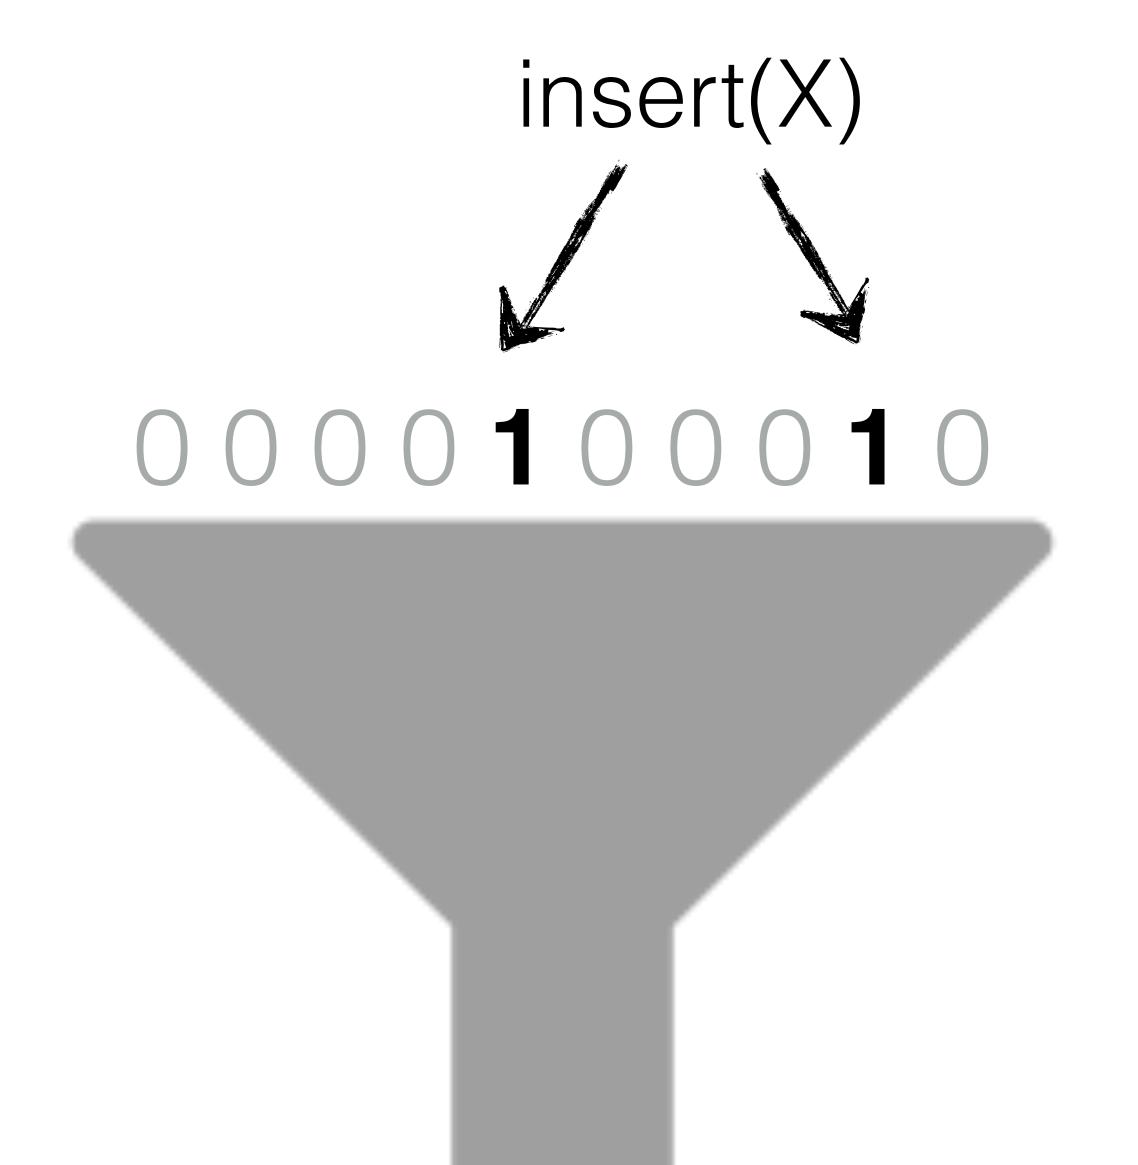
## Bloom Filters

Space/time Trade-Offs in Hash Coding with Allowable Errors Burton Howard Bloom. Communications of the ACM, 1970.

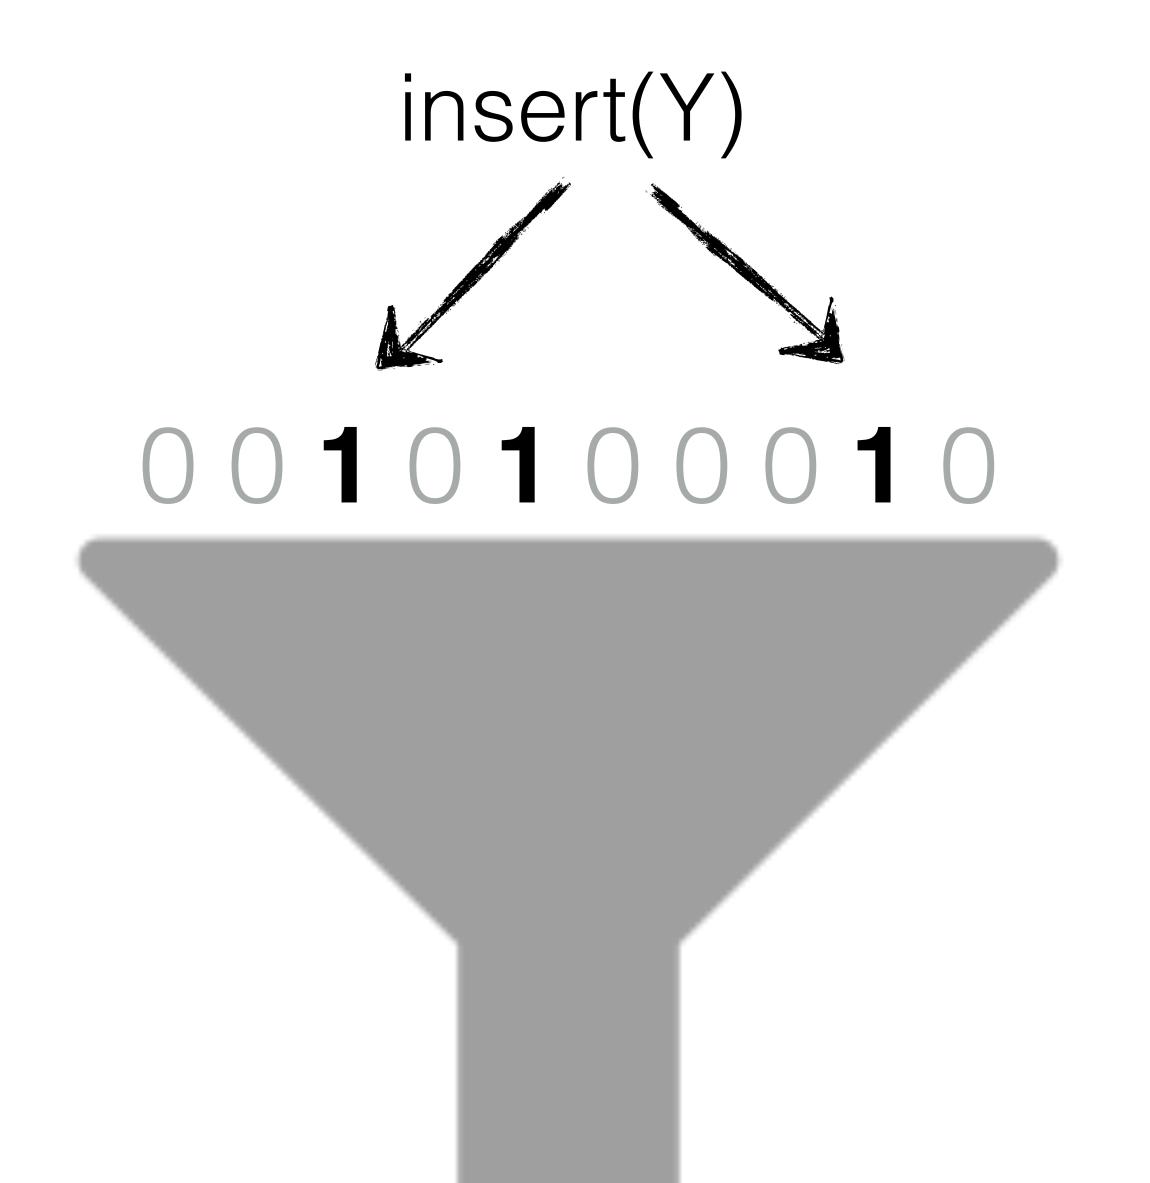
## k hash functions

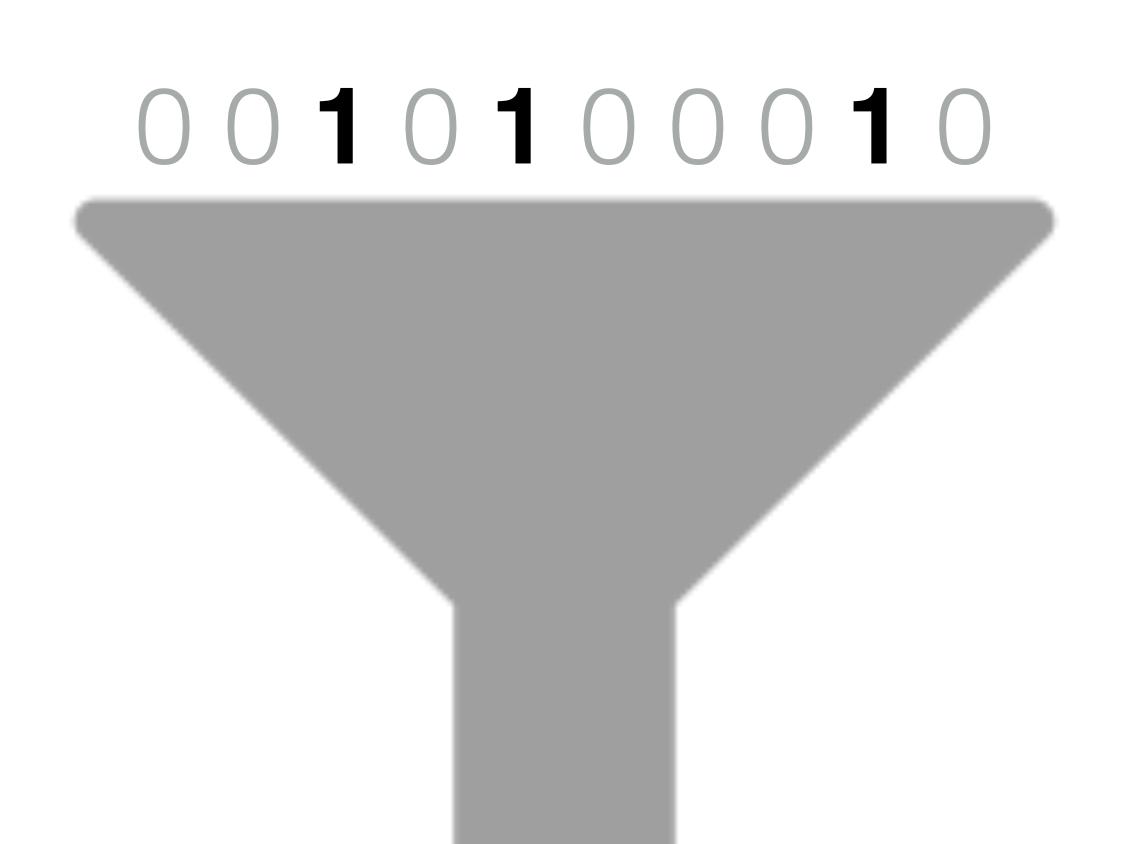


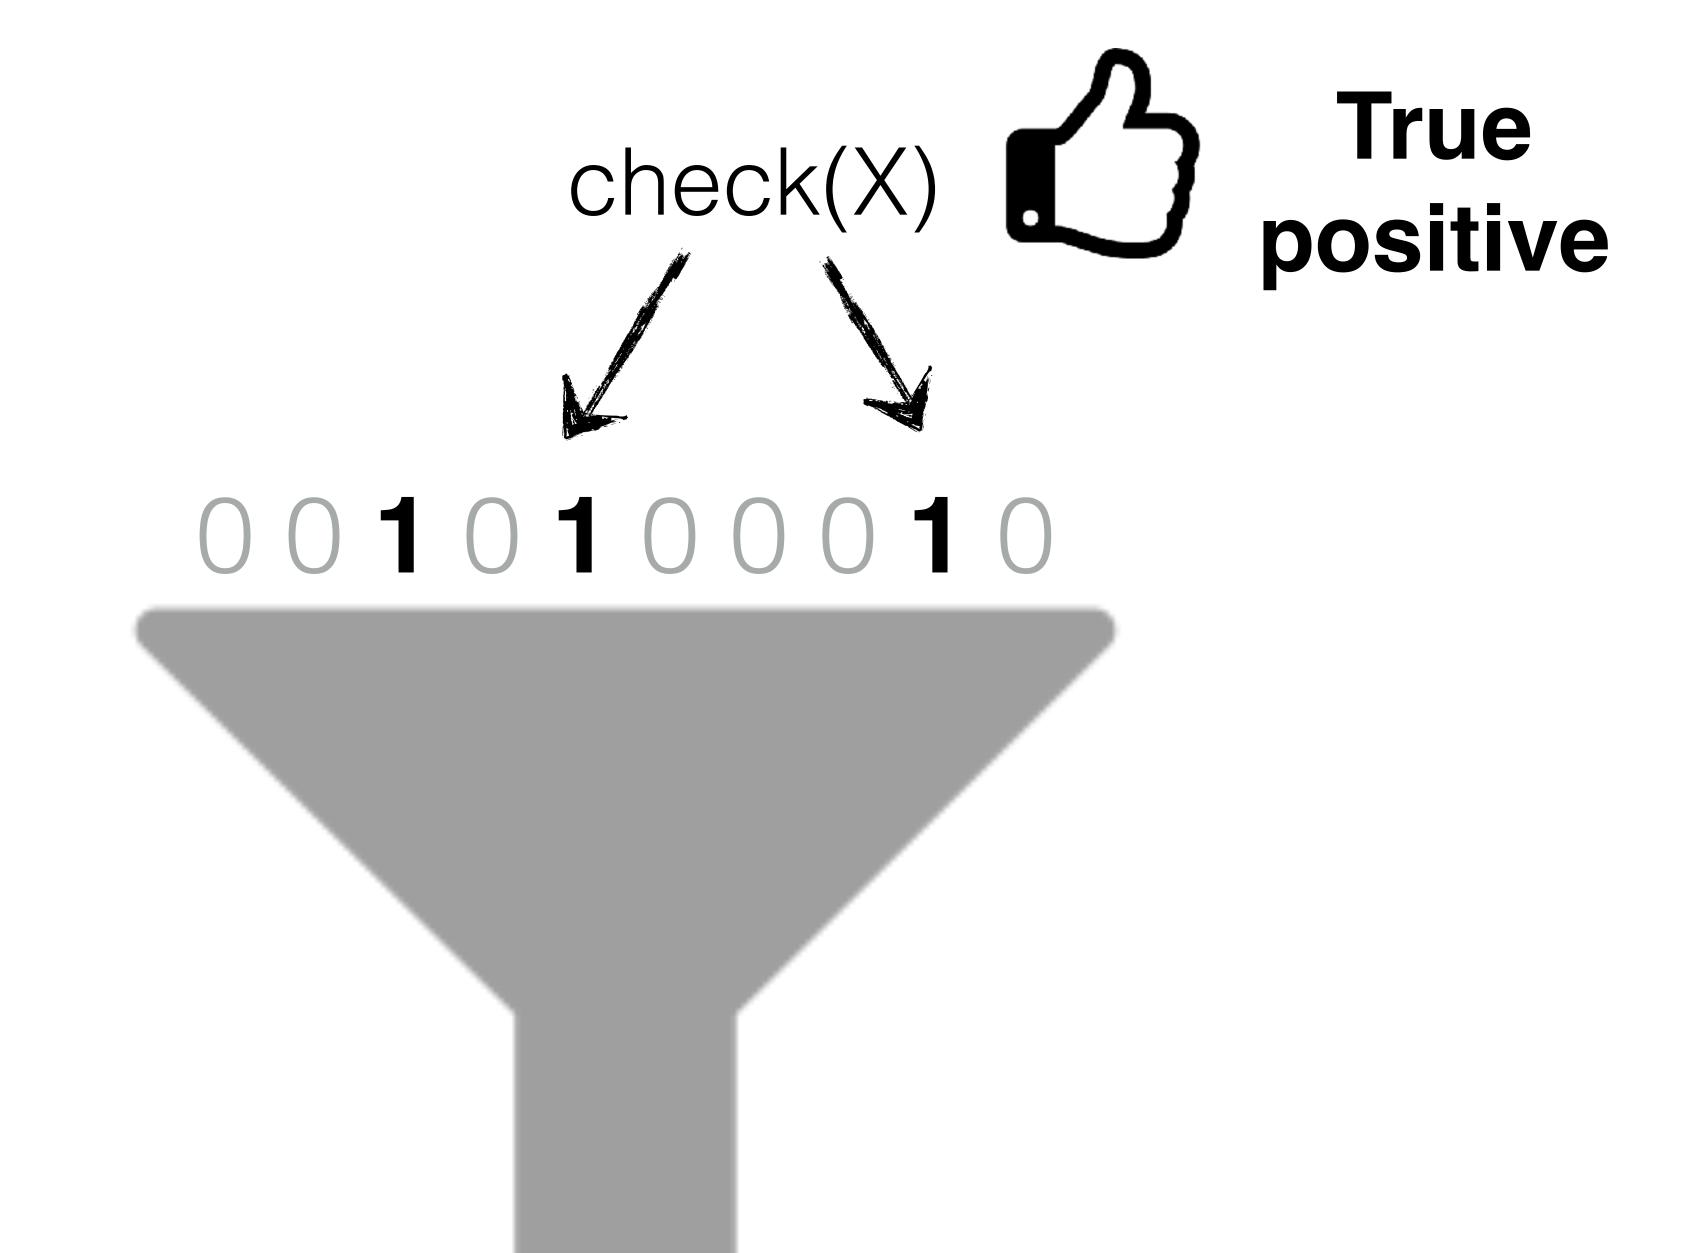
# insert: Set from 0 to 1 or keep 1

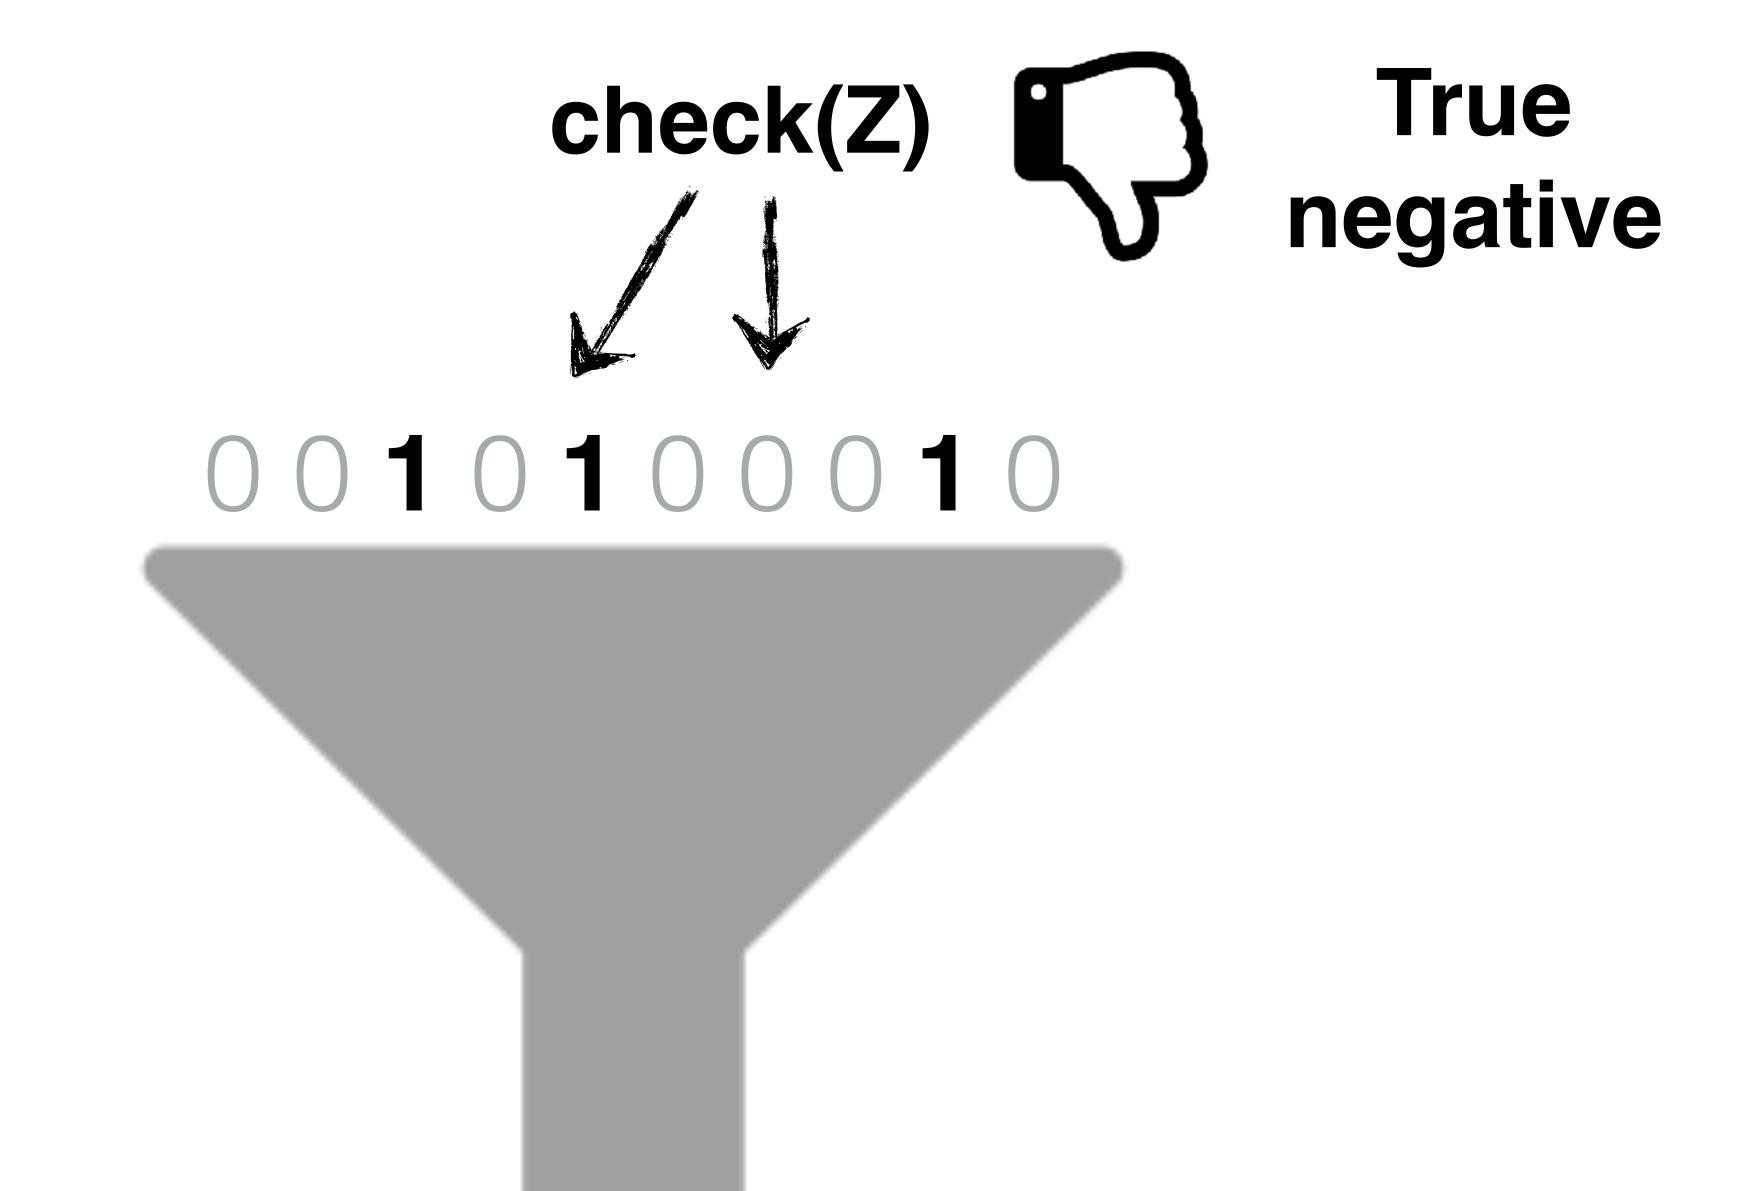


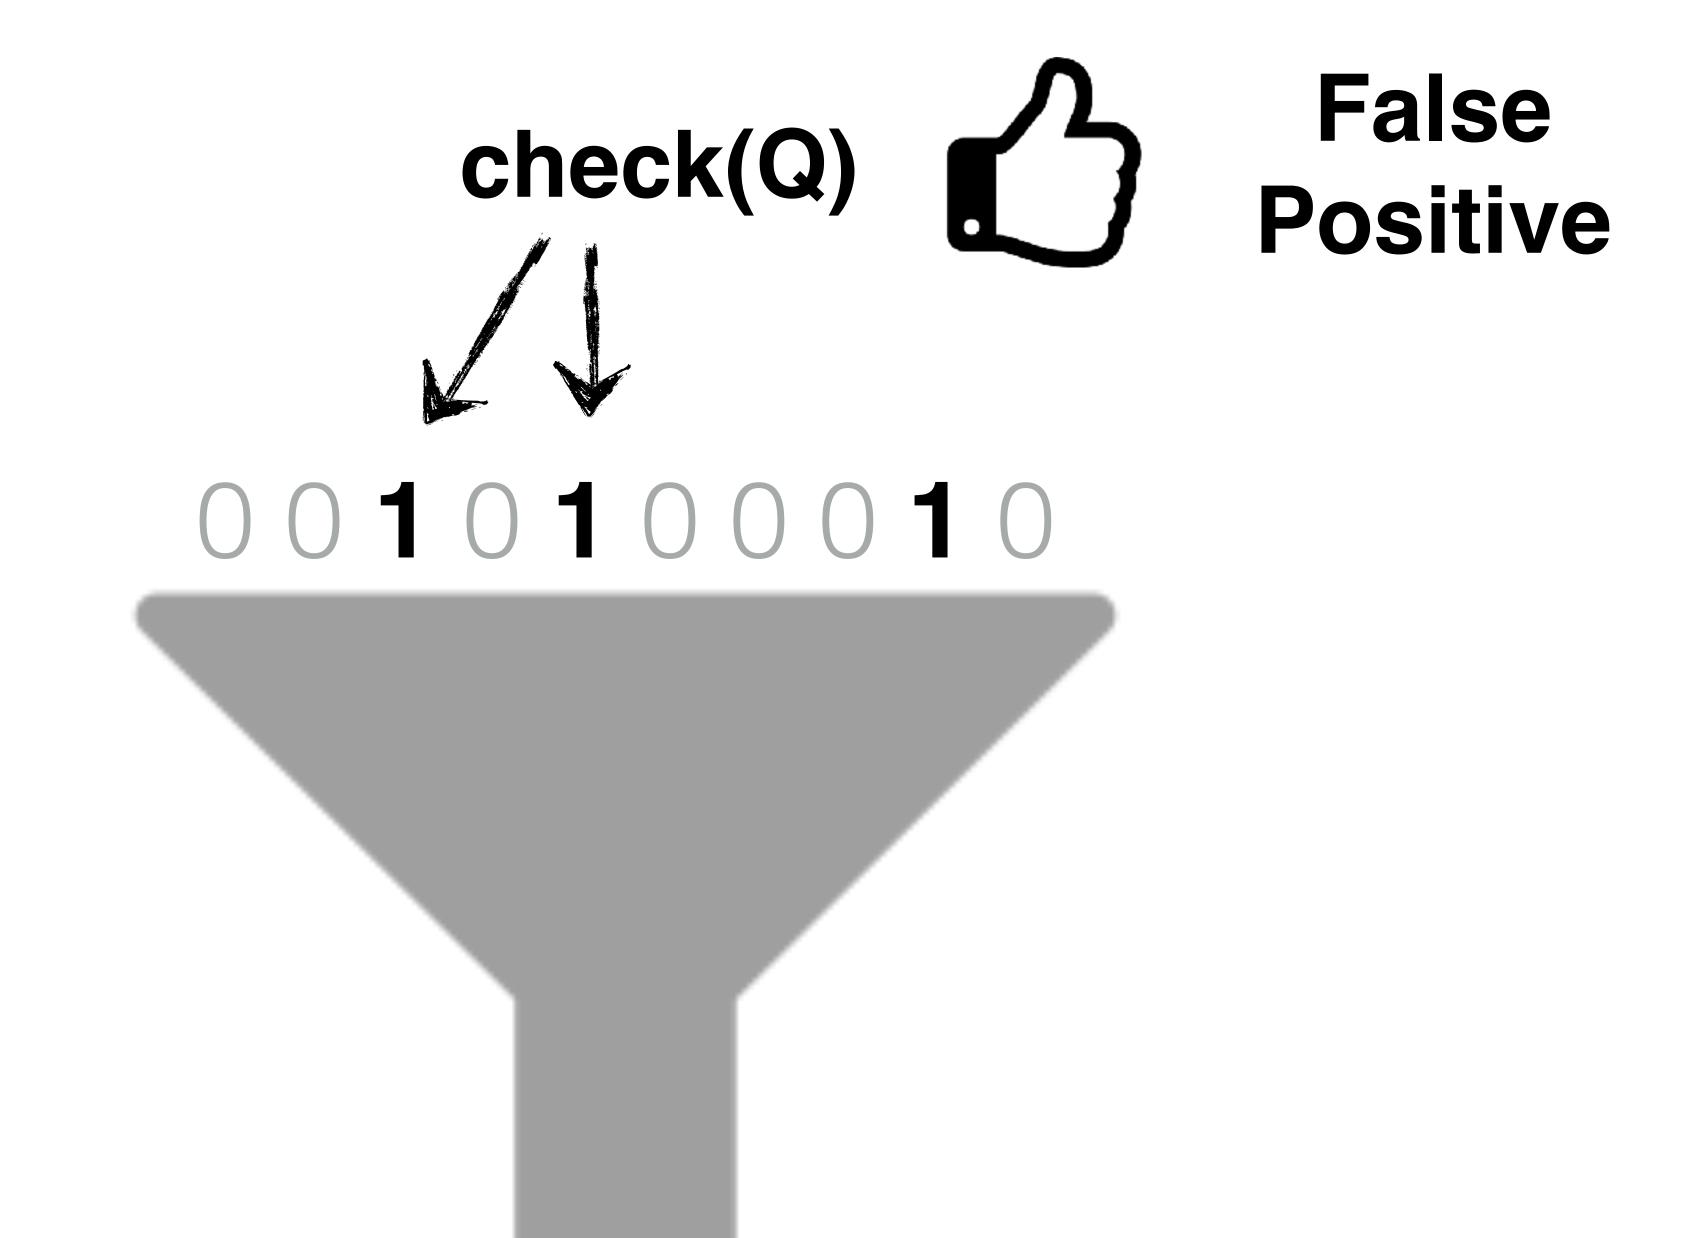
# insert: Set from 0 to 1 or keep 1



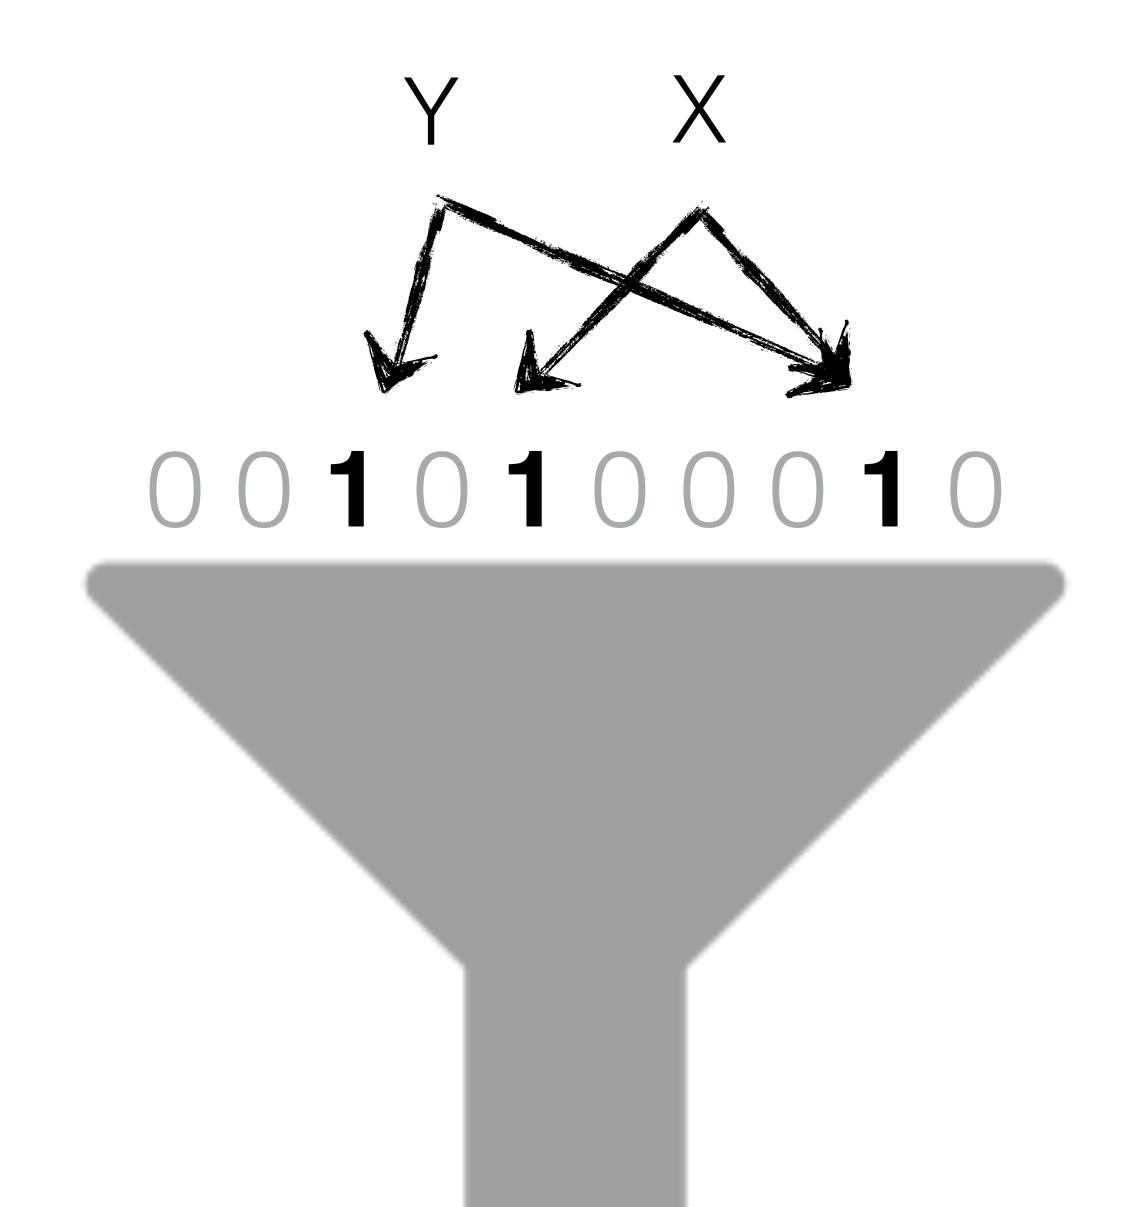




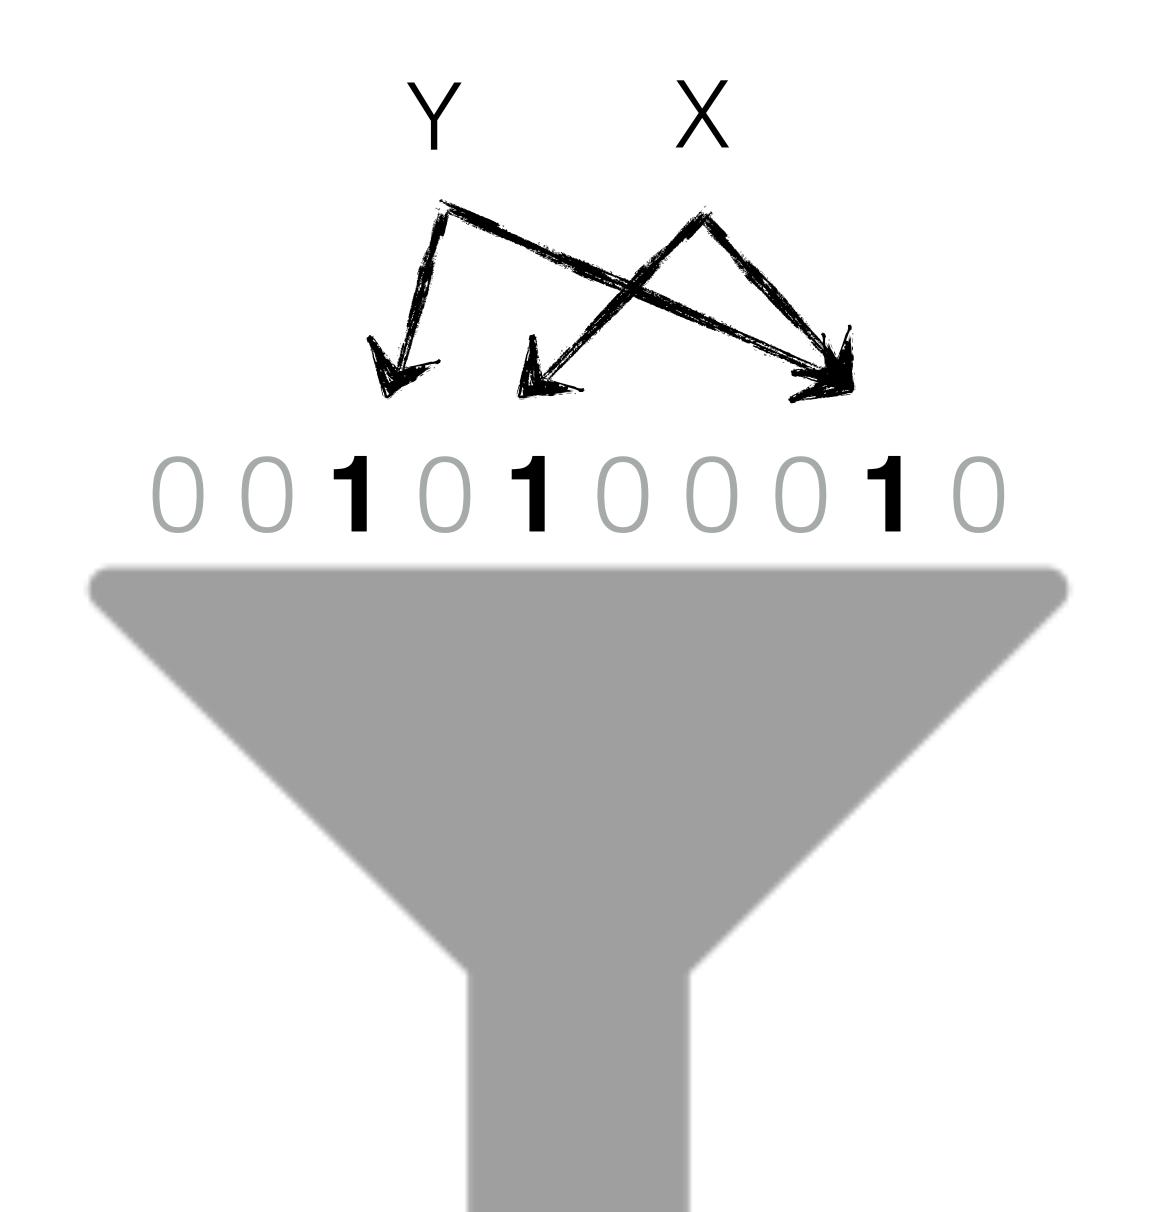




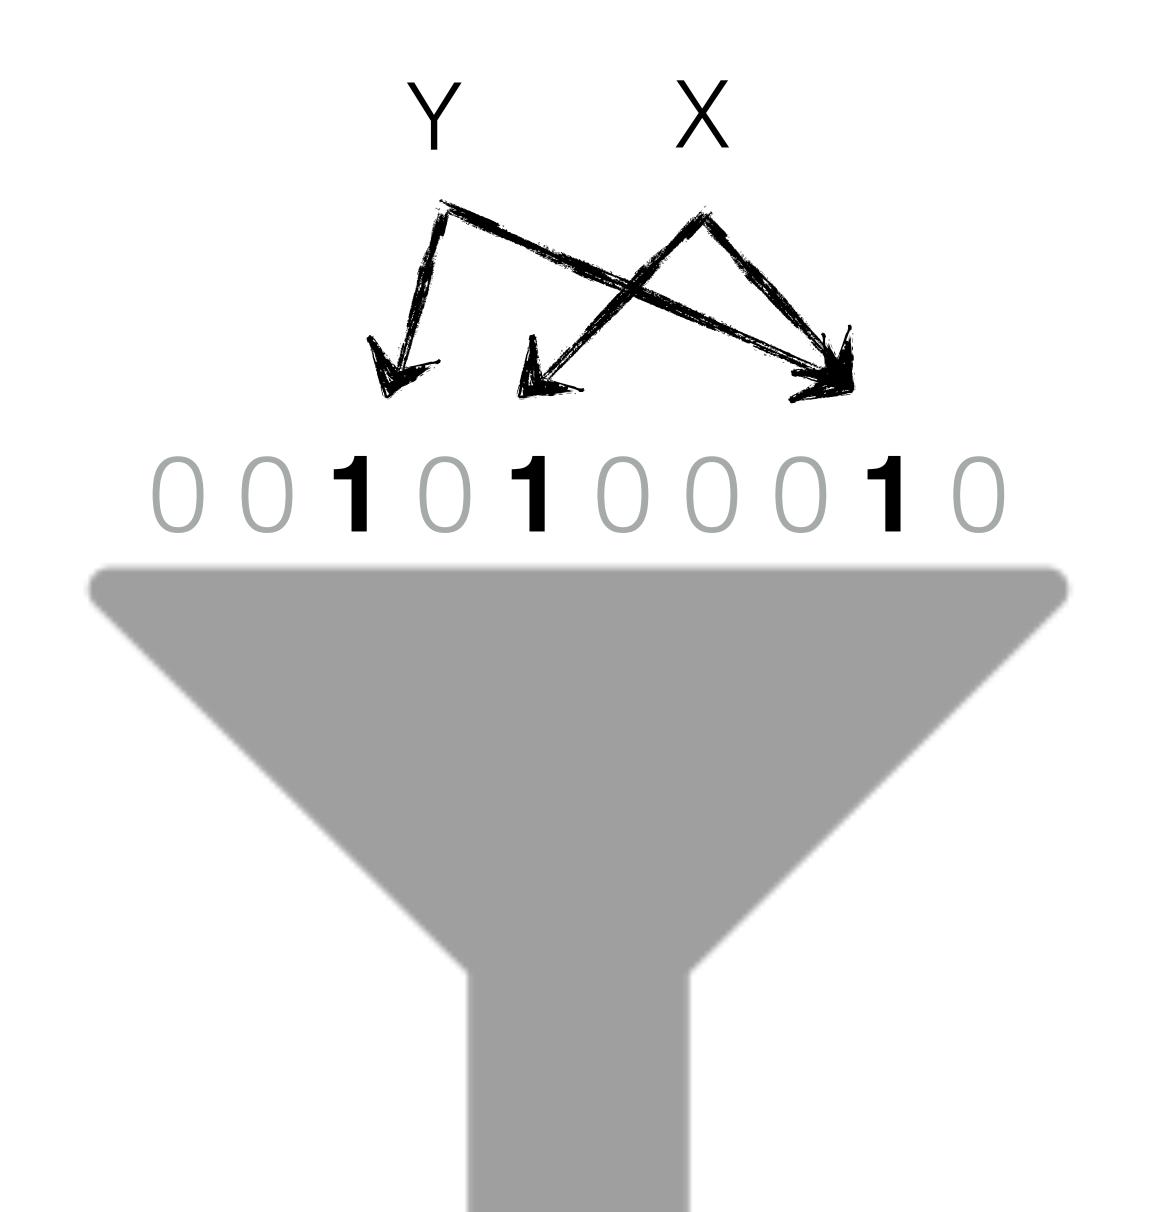
# No deletes - can lead to false negatives



# Thus, we consider it static



# Thus, we consider it static



## Construction contract



#### Construction contract

# Know specs in advance:

- N # entries to insert
- ε desired FPR



### Construction contract

Know specs in advance:

N - # entries to insert

ε - desired FPR

Allocate filter with:  $N \cdot ln(2) \cdot log_2(1/\epsilon)$  bits



### Construction contract

Know specs in advance:

N - # entries to insert

ε - desired FPR

Allocate filter with:  $N \cdot ln(2) \cdot log_2(1/\epsilon)$  bits

Insert N elements



### Construction contract

Know specs in advance:

N - # entries to insert

ε - desired FPR

Allocate filter with:  $N \cdot ln(2) \cdot log_2(1/\epsilon)$  bits

Insert N elements

Guarantee FPR of  $\epsilon$ 

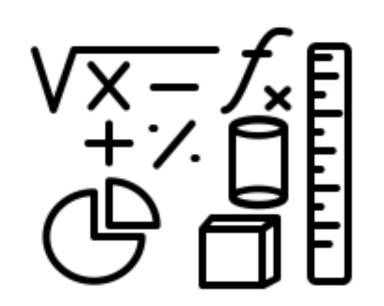


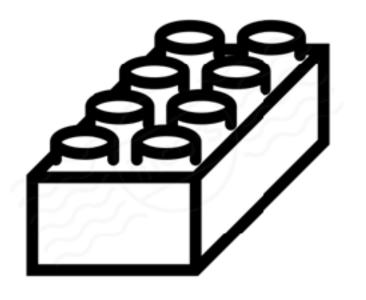
### **Bloom Filters**

Analysis

Blocking

Sectorization

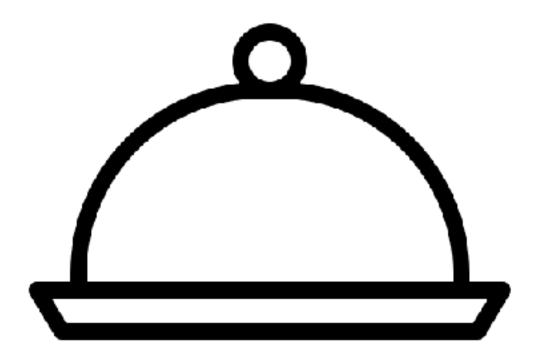




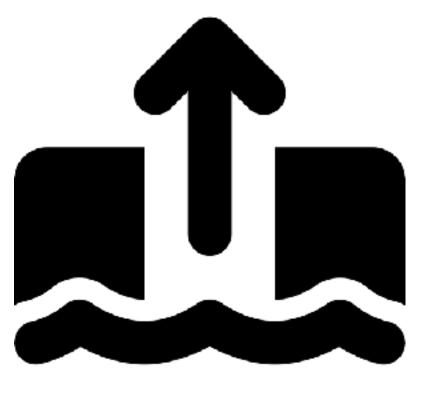


## Analysis

In CSC443



Now: ground up



Network Applications of Bloom Filters: A Survey

Andrei Broder and Michael Mitzenmacher.

Allerton Conference, 2002

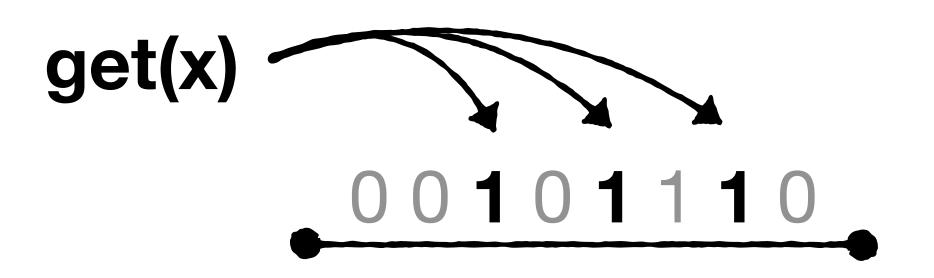
N: Total number of keys

K: # hash functions

N: Total number of keys

K: # hash functions

#### Probability that all k bits for a non-existing key are set?

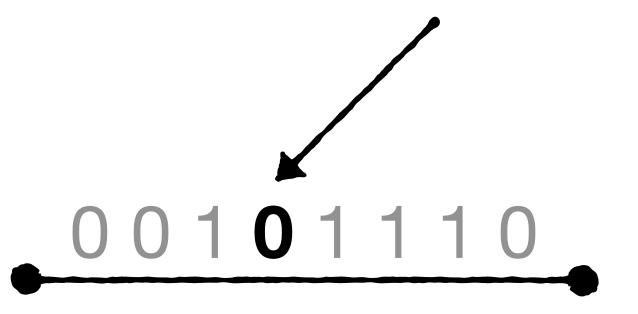


N: Total number of keys

K: # hash functions

Probability that all k bits for a non-existing key are set?

Probability that some random bit is still not set after N insertions?



M: Total number of bits

N: Total number of keys

K: # hash functions

Probability that all k bits for a non-existing key are set?

1

Probability that some random bit is still not set after N insertions?

1

Probability that some random bit is still not set after 1 insertion?

M: Total number of bits

N: Total number of keys

K: # hash functions

Probability that all k bits for a non-existing key are set?



Probability that some random bit is still not set after N insertions?



Probability that some random bit is still not set after 1 insertion?



Probability that some random bit is set after 1 hash function?

M: Total number of bits

N: Total number of keys

K: # hash functions

Probability that all k bits for a non-existing key are set?

Probability that some random bit is still not set after N insertions?



Probability that some random bit is still not set after 1 insertion?



Probability that some random bit is **not** set after 1 hash function?

M: Total number of bits

N: Total number of keys

K: # hash functions

Probability that all k bits for a non-existing key are set?



Probability that some random bit is still not set after N insertions?



Probability that some random bit is still not set after 1 insertion?

$$(1-1/M)^{K}$$

M: Total number of bits

N: Total number of keys

K: # hash functions

Probability that all k bits for a non-existing key are set?



Probability that some random bit is still not set after N insertions?

$$(1-1/M)^{KN}$$

M: Total number of bits

N: Total number of keys

K: # hash functions

Probability that all k bits for a non-existing key are set?



Probability that some random bit is still not set after N insertions?

$$(1-1/M)^{KN}$$

Known identity: 
$$(1-1/M)^{M} = e^{-1}$$

For any M

M: Total number of bits

N: Total number of keys

K: # hash functions

Probability that all k bits for a non-existing key are set?



Probability that some random bit is still not set after N insertions?

$$((1-1/M)M)KN/M$$

Known identity: 
$$(1-1/M)^{M} = e^{-1}$$

For any M

M: Total number of bits

N: Total number of keys

K: # hash functions

Probability that all k bits for a non-existing key are set?



Probability that some random bit is still not set after N insertions?

$$(e^{-1})$$
 KN/M

M: Total number of bits

N: Total number of keys

K: # hash functions

Probability that all k bits for a non-existing key are set?



Probability that some random bit is still not set after N insertions?

e-KN/M

M: Total number of bits

N: Total number of keys

K: # hash functions

Probability that all k bits for a non-existing key are set?



Probability that some random bit is still not set after N insertions?

M: Total number of bits

N: Total number of keys

K: # hash functions

Probability that all k bits for a non-existing key are set?



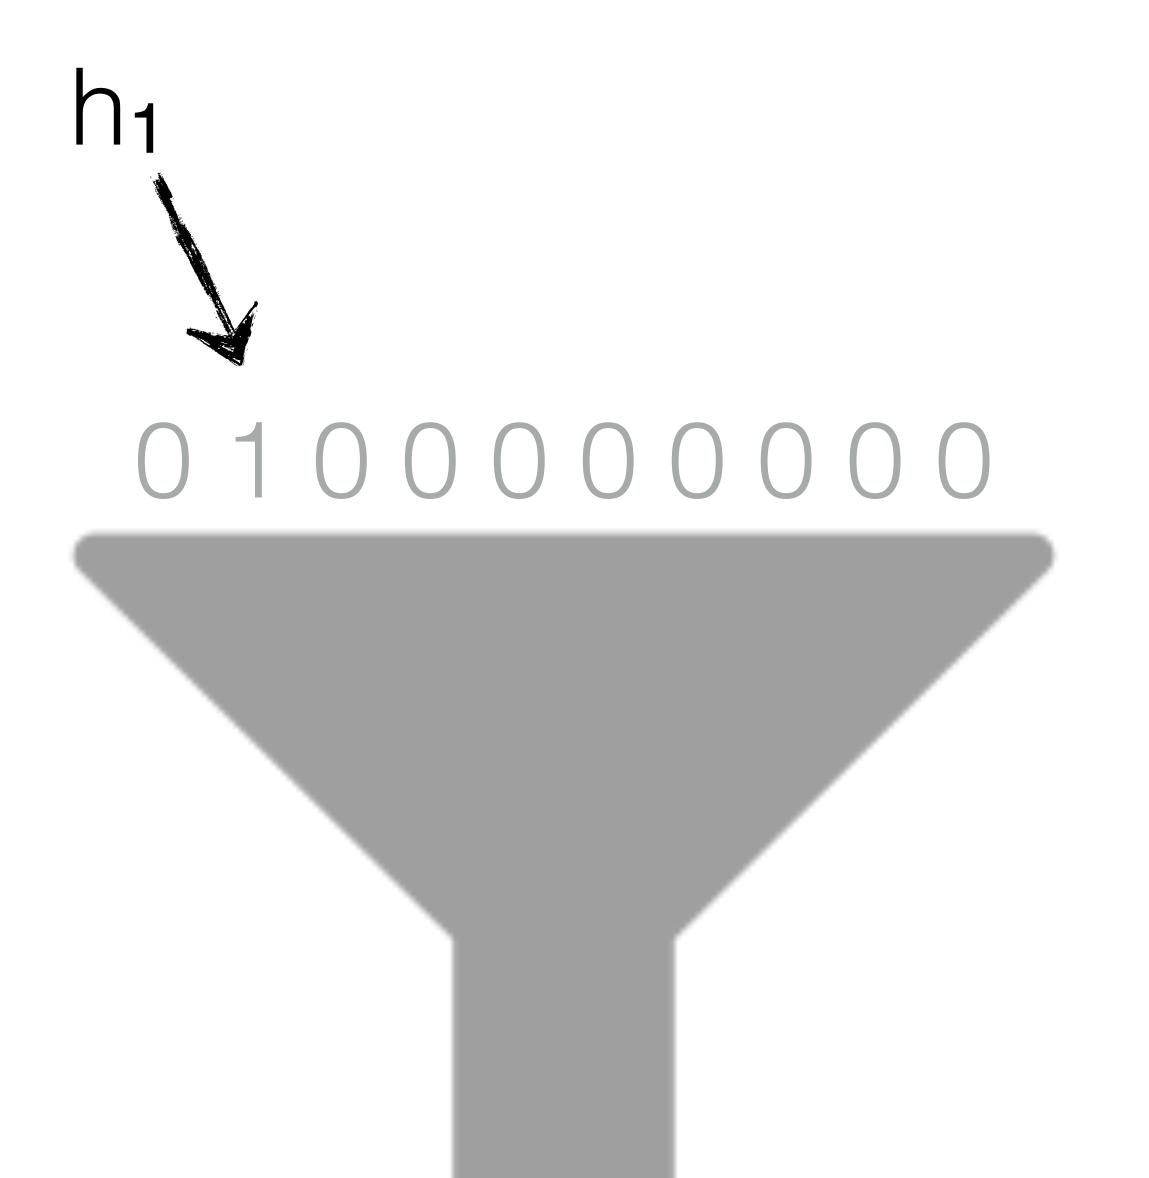
Probability that some random bit is set after N insertions?

N: Total number of keys

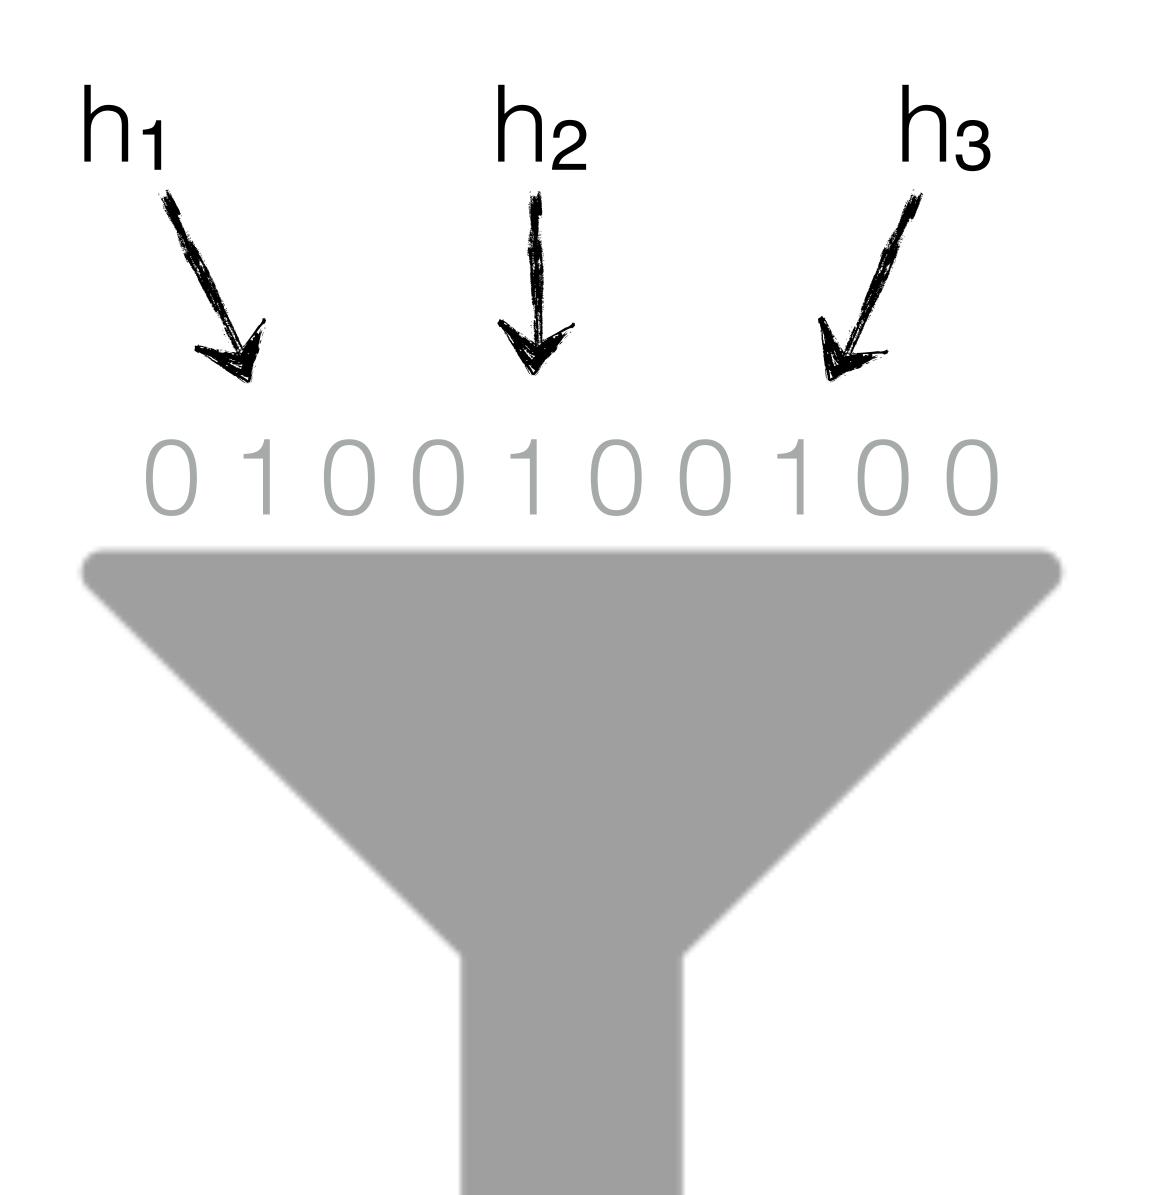
K: # hash functions

Probability that all k bits for a non-existing key are set?

 $(1-e^{-KN/M})K$ 

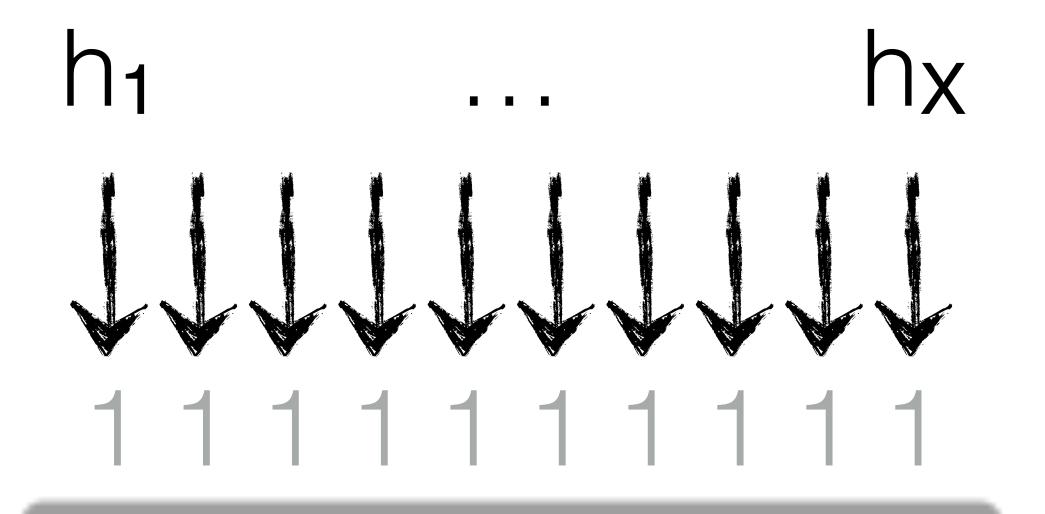


One is too few: false positive occurs whenever we hit a 1



One is too few: false positive occurs whenever we hit a 1

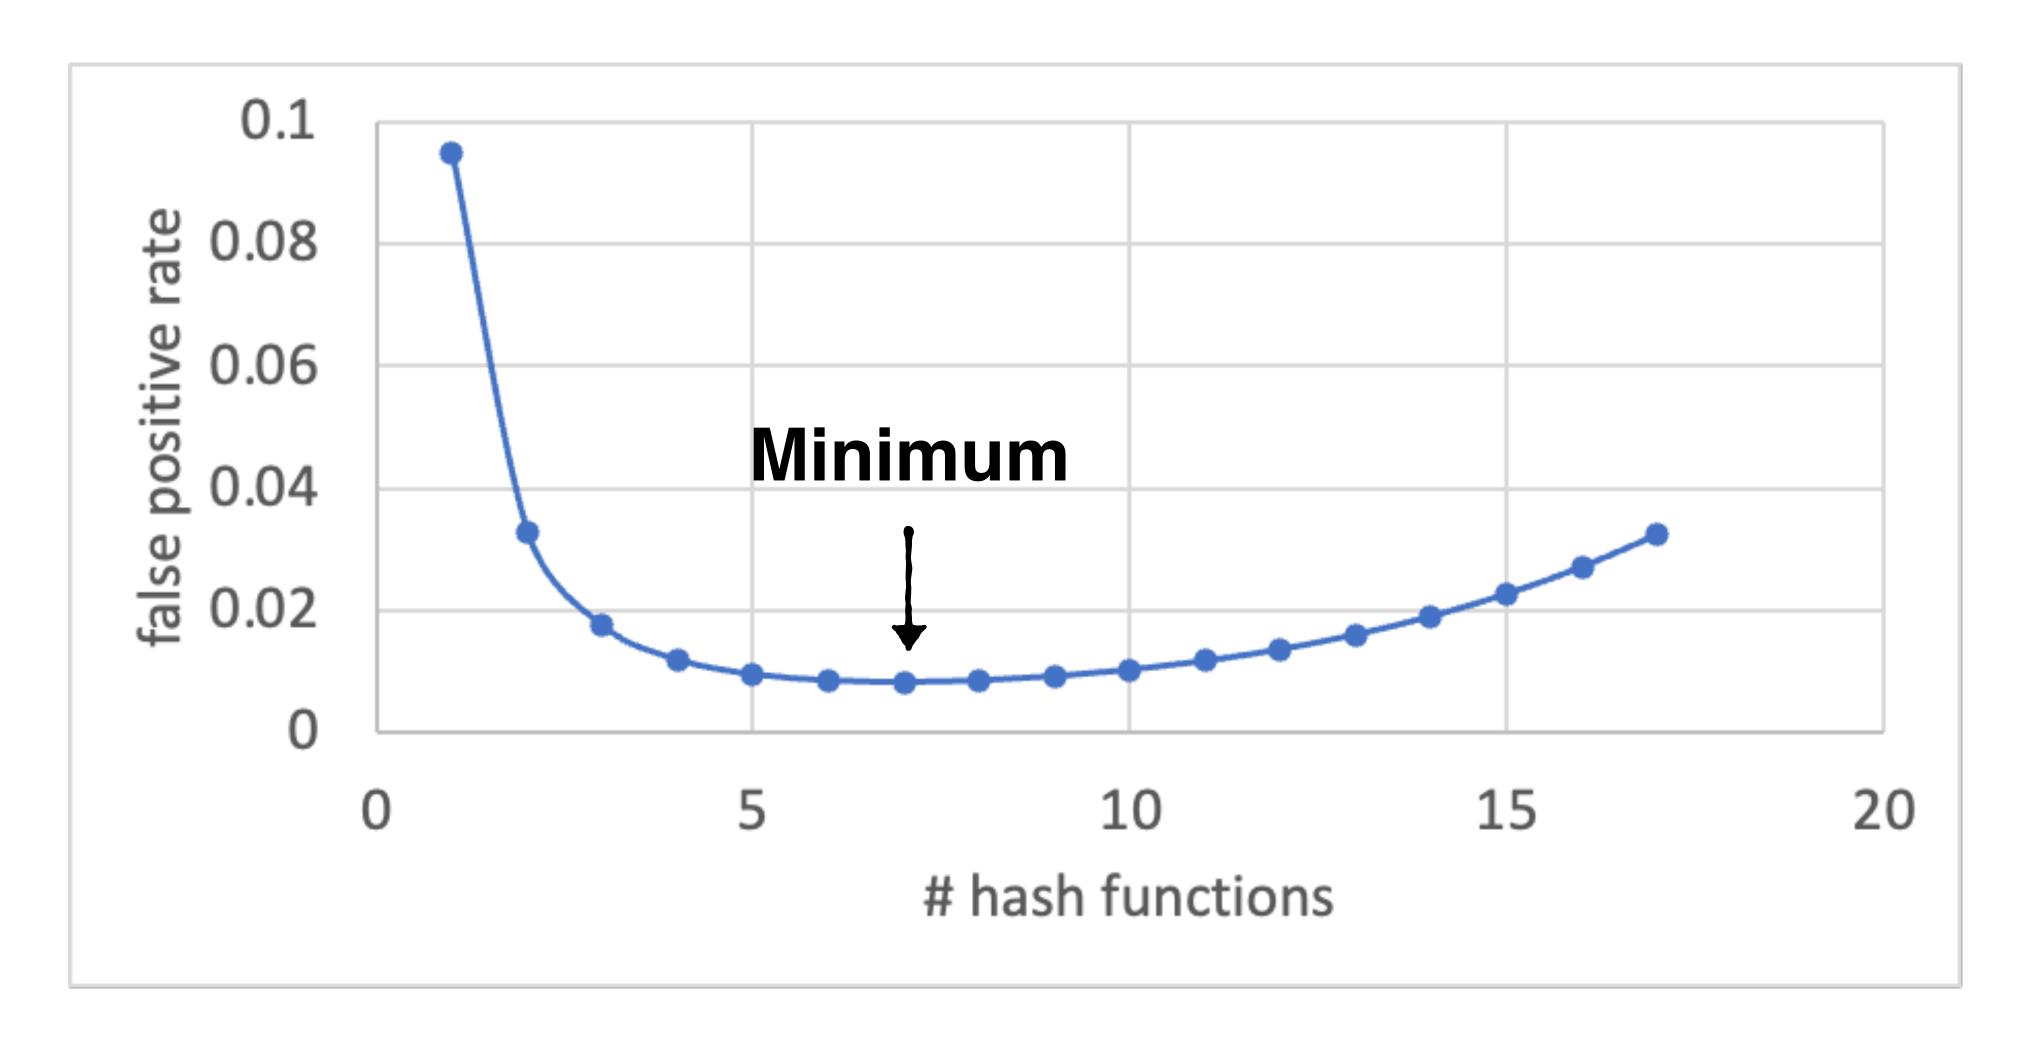
By adding hash functions, we initially decrease the false positive rate (FPR).



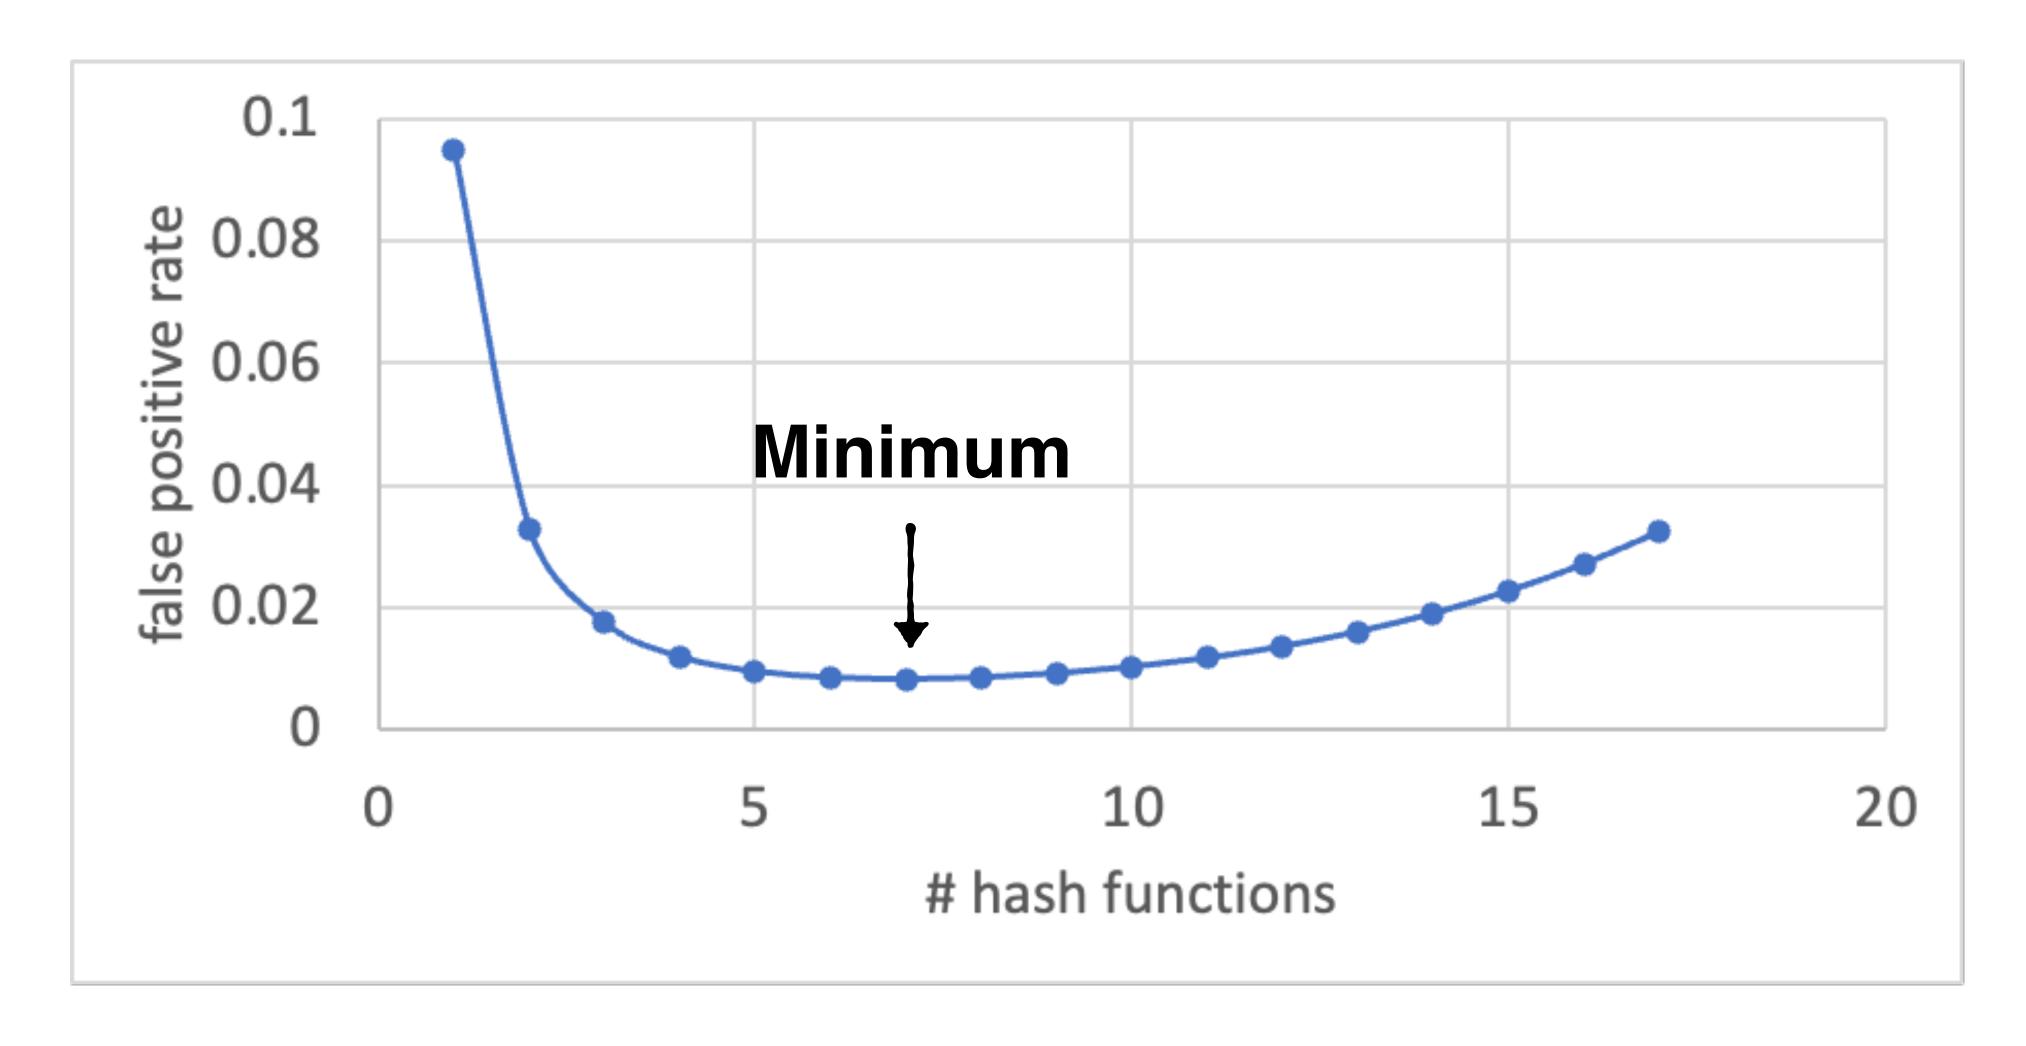
One is too few: false positive occurs whenever we hit a 1

By adding hash functions, we initially decrease the false positive rate (FPR).

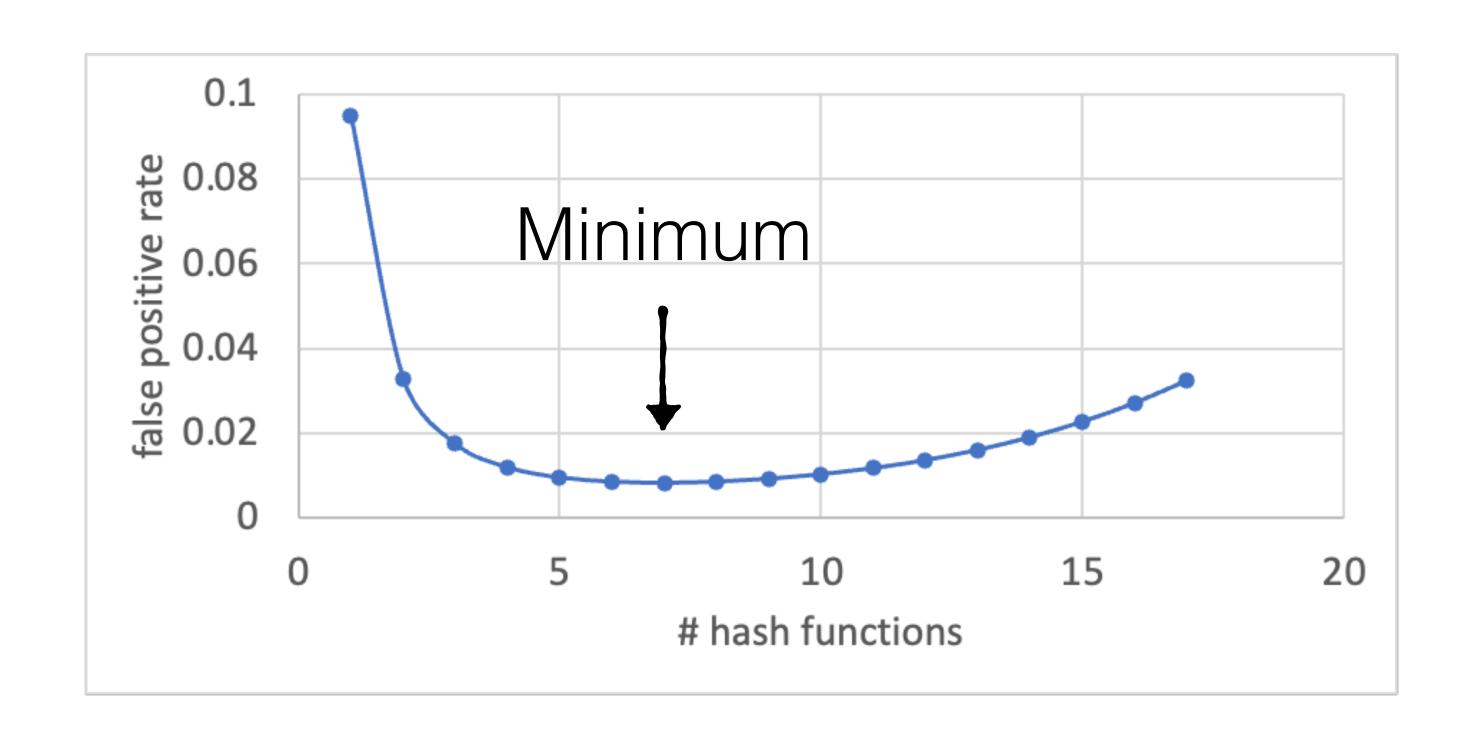
But too many hash functions wind up increasing the FPR.



(Drawn for a filter using 10 bits per entry)



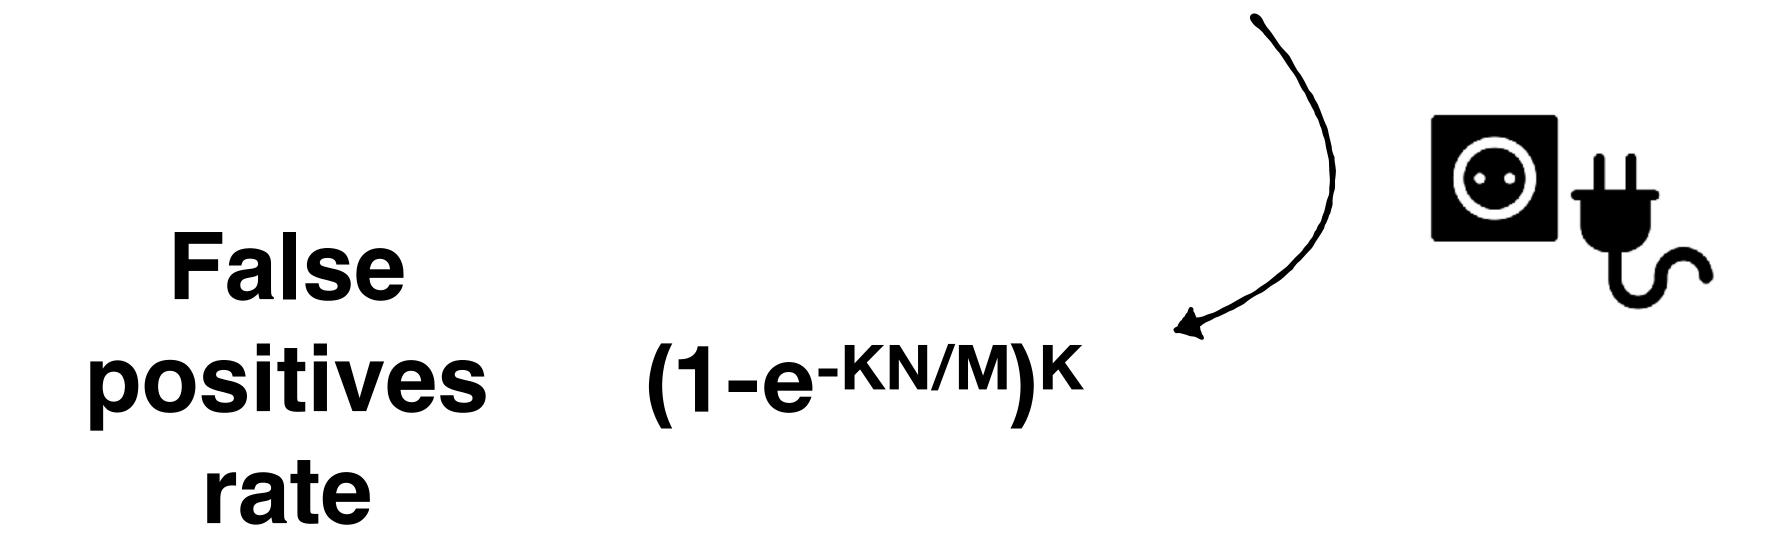
Differentiate (1-e-KN/M)K with respect to K



Differentiate (1-e-KN/M)K with respect to K

Optimal # hash functions  $k = ln(2) \cdot M/N$ 

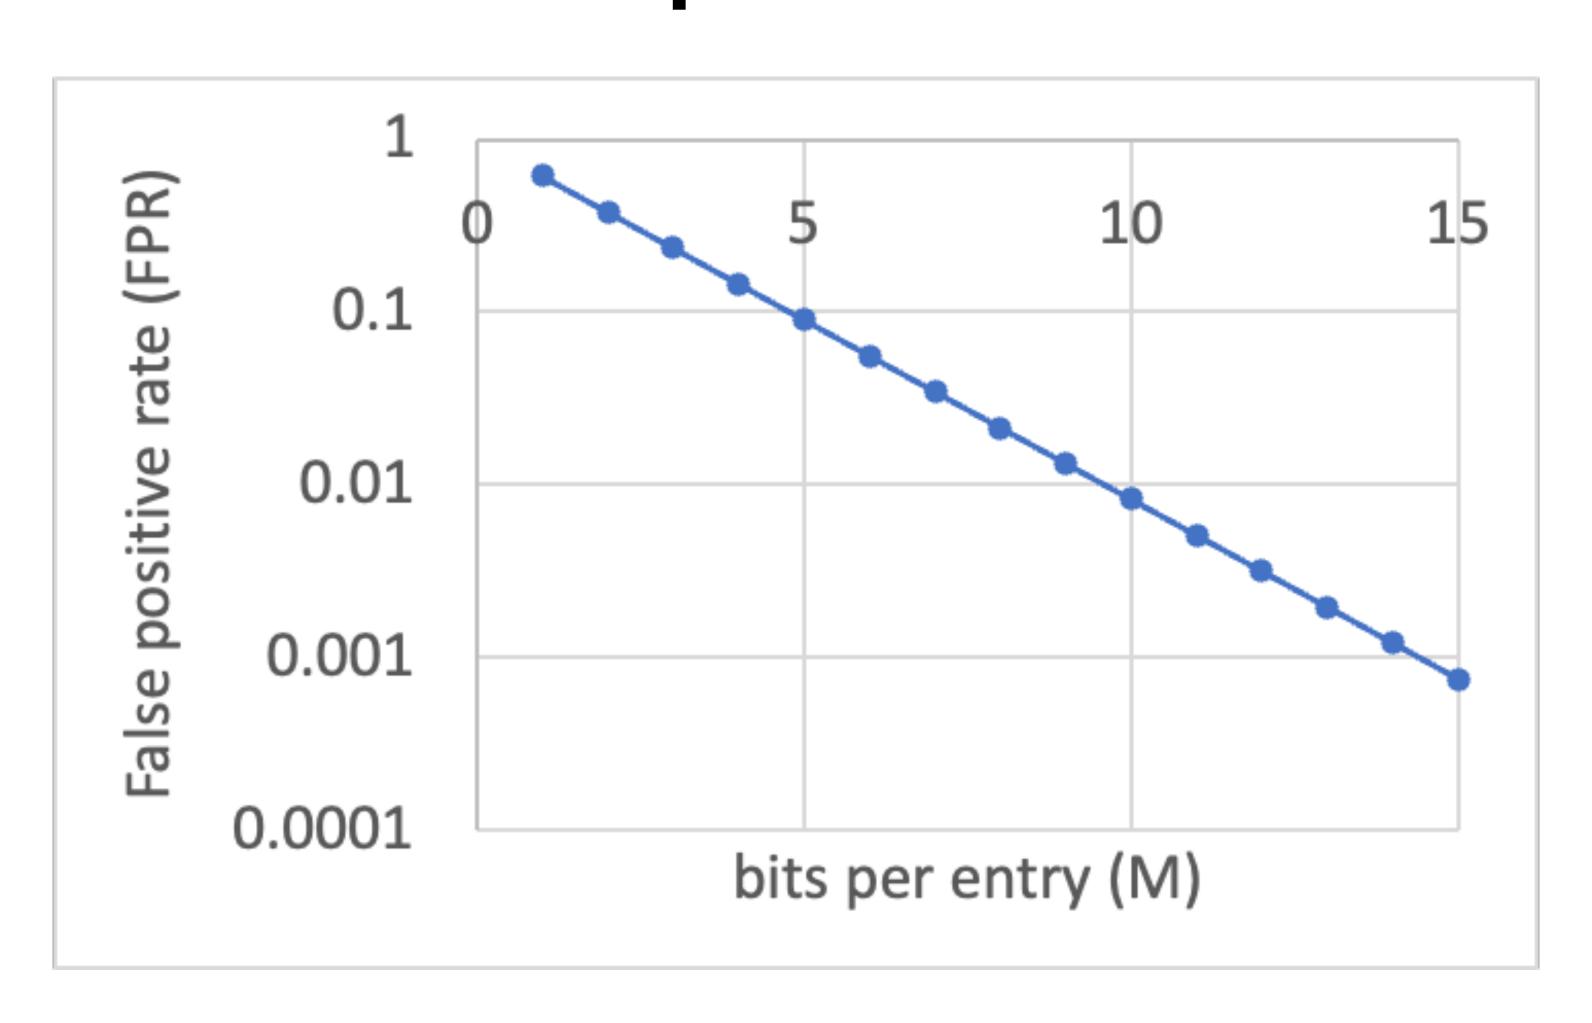
(e.g. with Wolfram Alpha)



False positives (1-e-KN/M)K rate

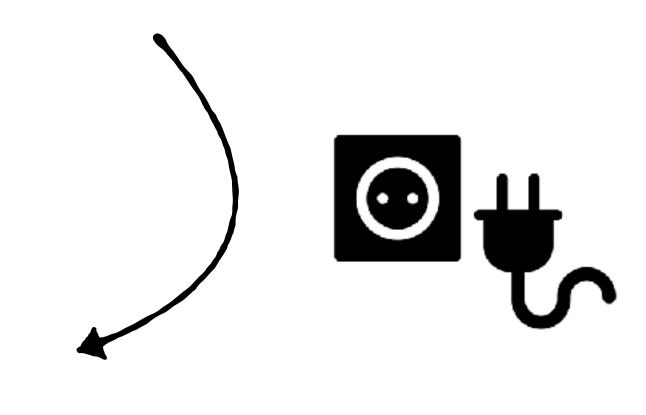
assuming the optimal # hash functions,

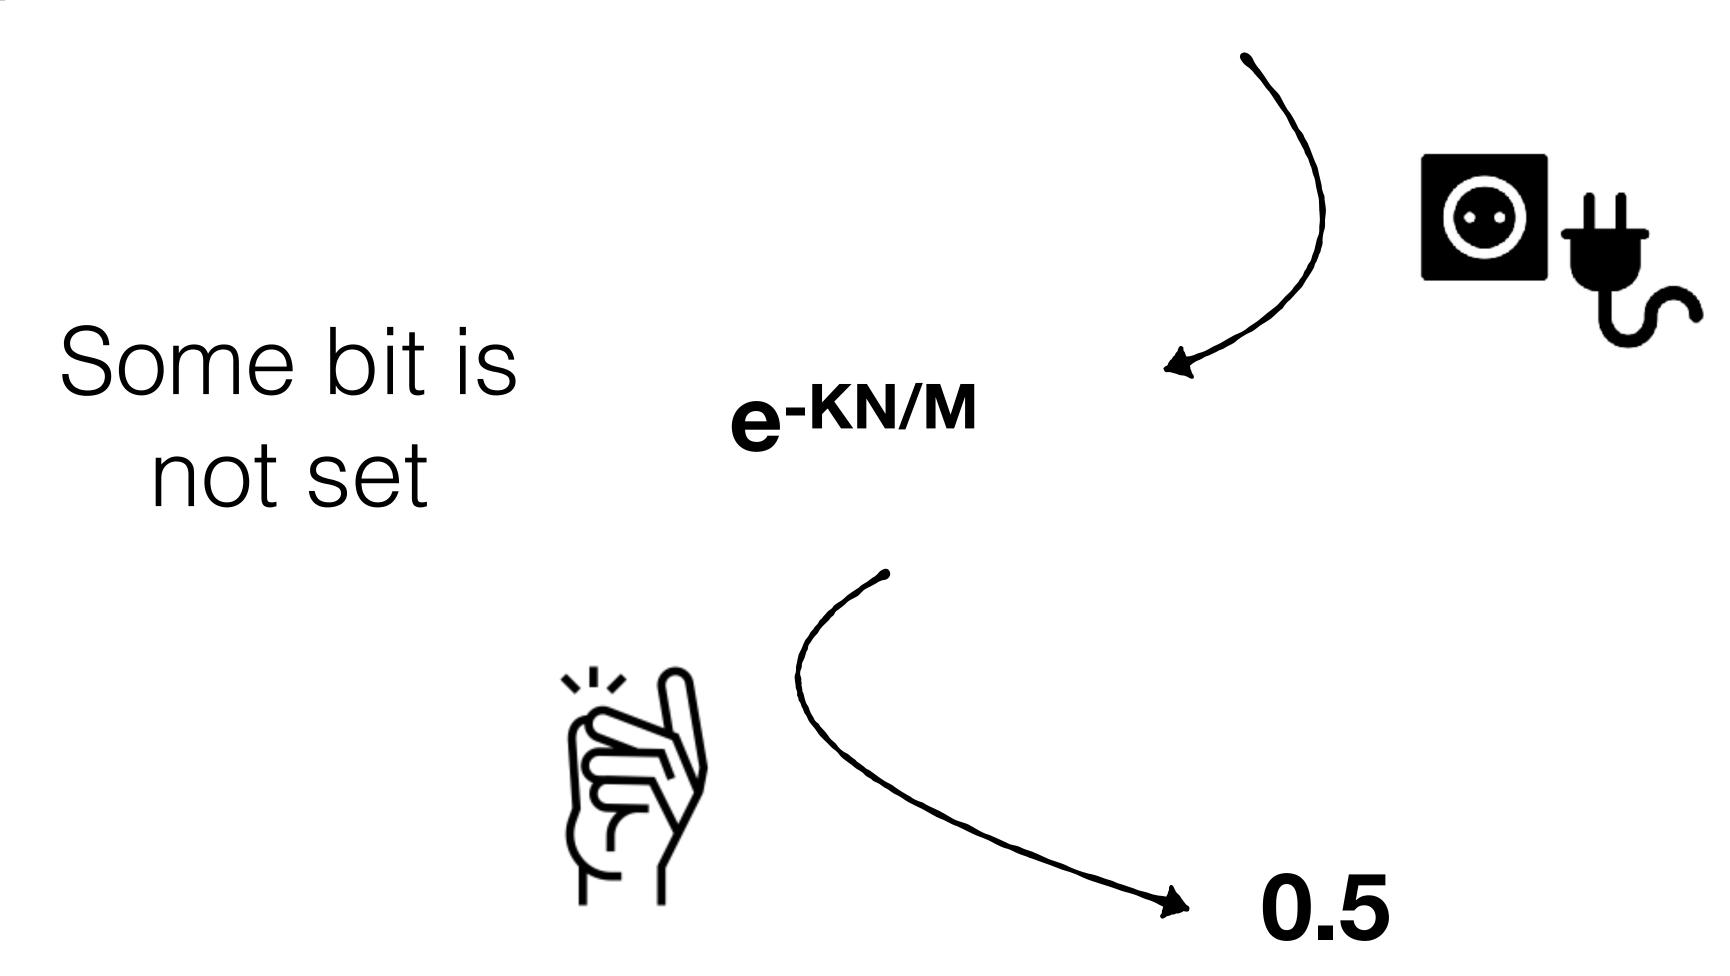
# false positive rate = $2-M/N \cdot ln(2)$

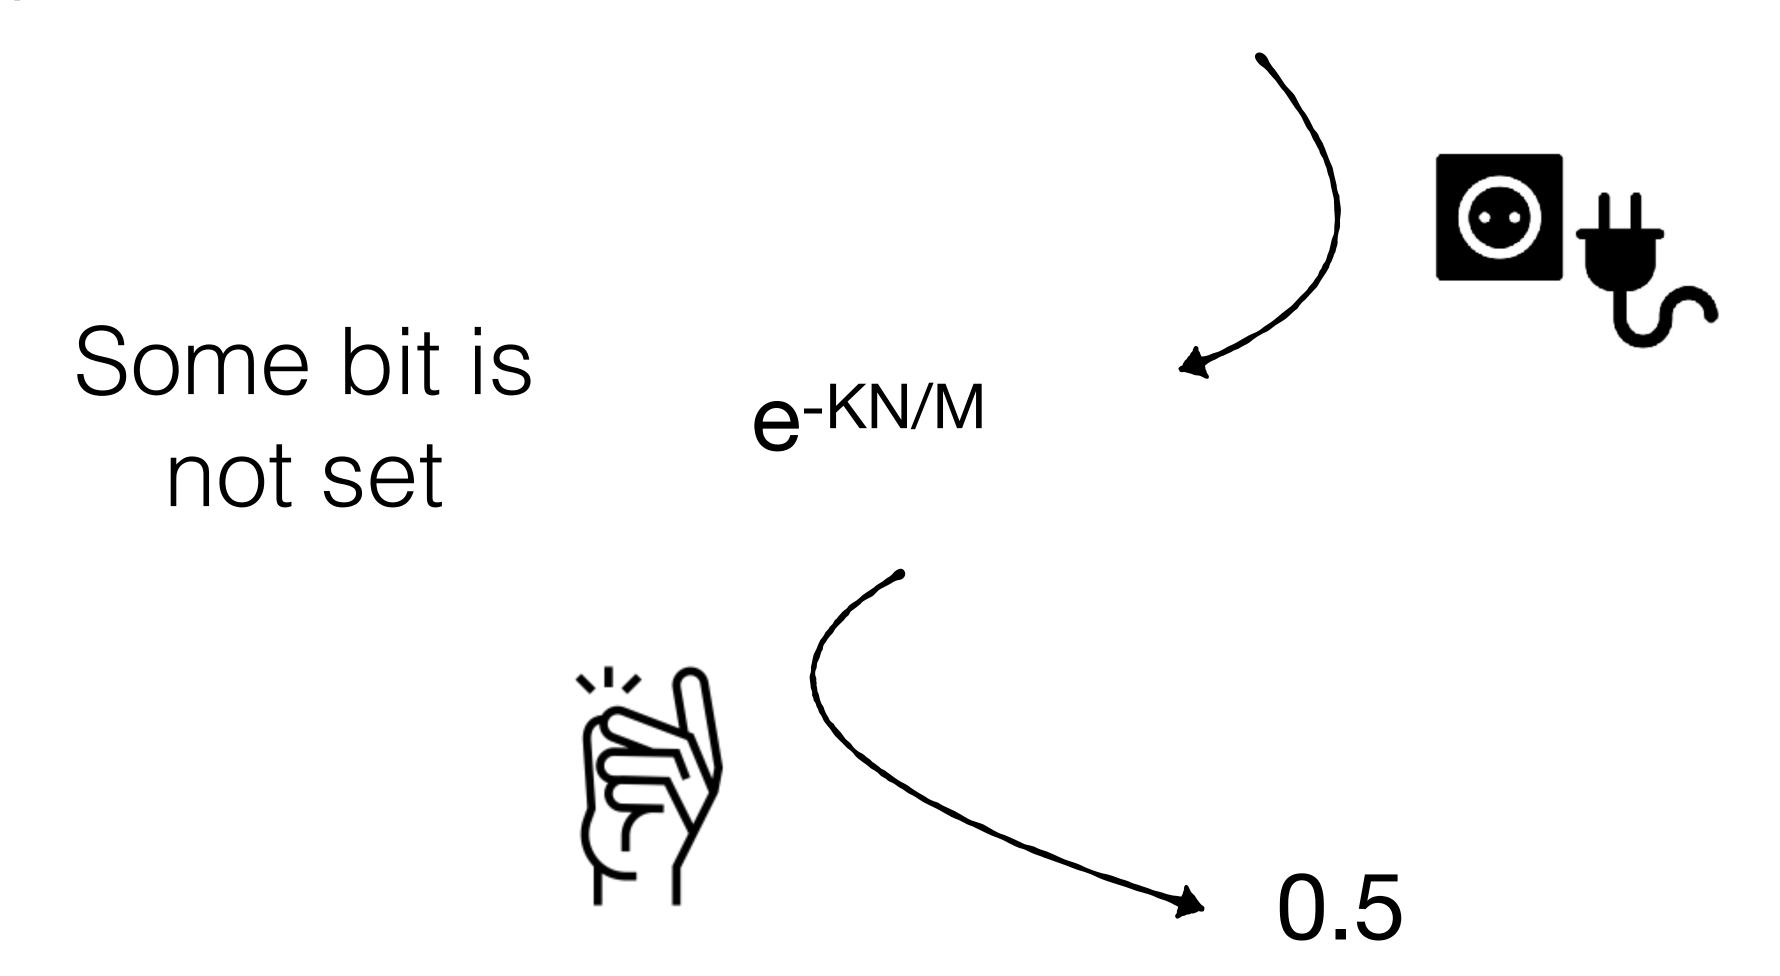


Some bit is not set

e-KN/M







50% of all bits are zero once the filter is full

### Operation Costs (in hash functions computed)



Insertion =



Positive Query =



Negative Query =

## Operation Costs (in hash functions computed)



Insertion = M/N·In(2)



Positive Query =



Negative Query =

## Operation Costs (in hash functions computed)



Insertion =  $M/N \cdot ln(2)$ 



 $\bigcirc$  Positive Query = M/N · In(2)



Negative Query =



Insertion =  $M/N \cdot ln(2)$ 



Solution Positive Query =  $M/N \cdot ln(2)$ 



Negative Query =

(50% of bits are zeros)



Insertion =  $M/N \cdot ln(2)$ 



Positive Query =  $M/N \cdot ln(2)$ 



**Avg. Negative Query =**  $1 + 1/2 (1 + 1/2 \cdot (...))$ 

(50% of bits are zeros)



Insertion =  $M/N \cdot ln(2)$ 



Positive Query =  $M/N \cdot ln(2)$ 



Avg. Negative Query = 1 + 1/2 + 1/4 + ... = 2

(50% of bits are zeros)



Insertion =  $M/N \cdot ln(2)$ 



Positive Query =  $M/N \cdot ln(2)$ 



Avg. Negative Query = 2

Full analysis from ground up:)



Insertion =  $M/N \cdot ln(2)$ 



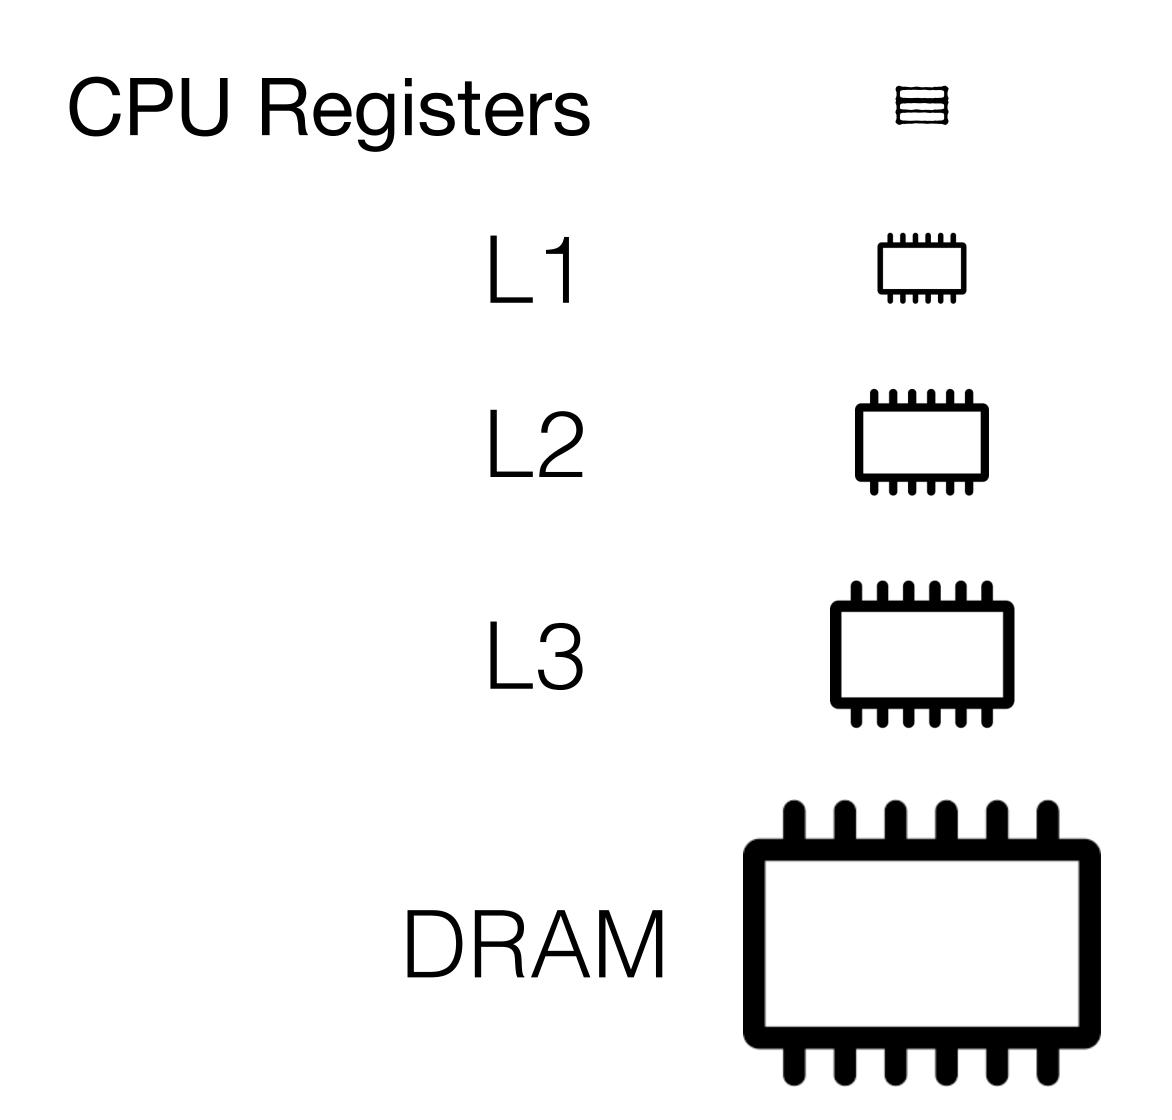
Positive Query =  $M/N \cdot ln(2)$ 

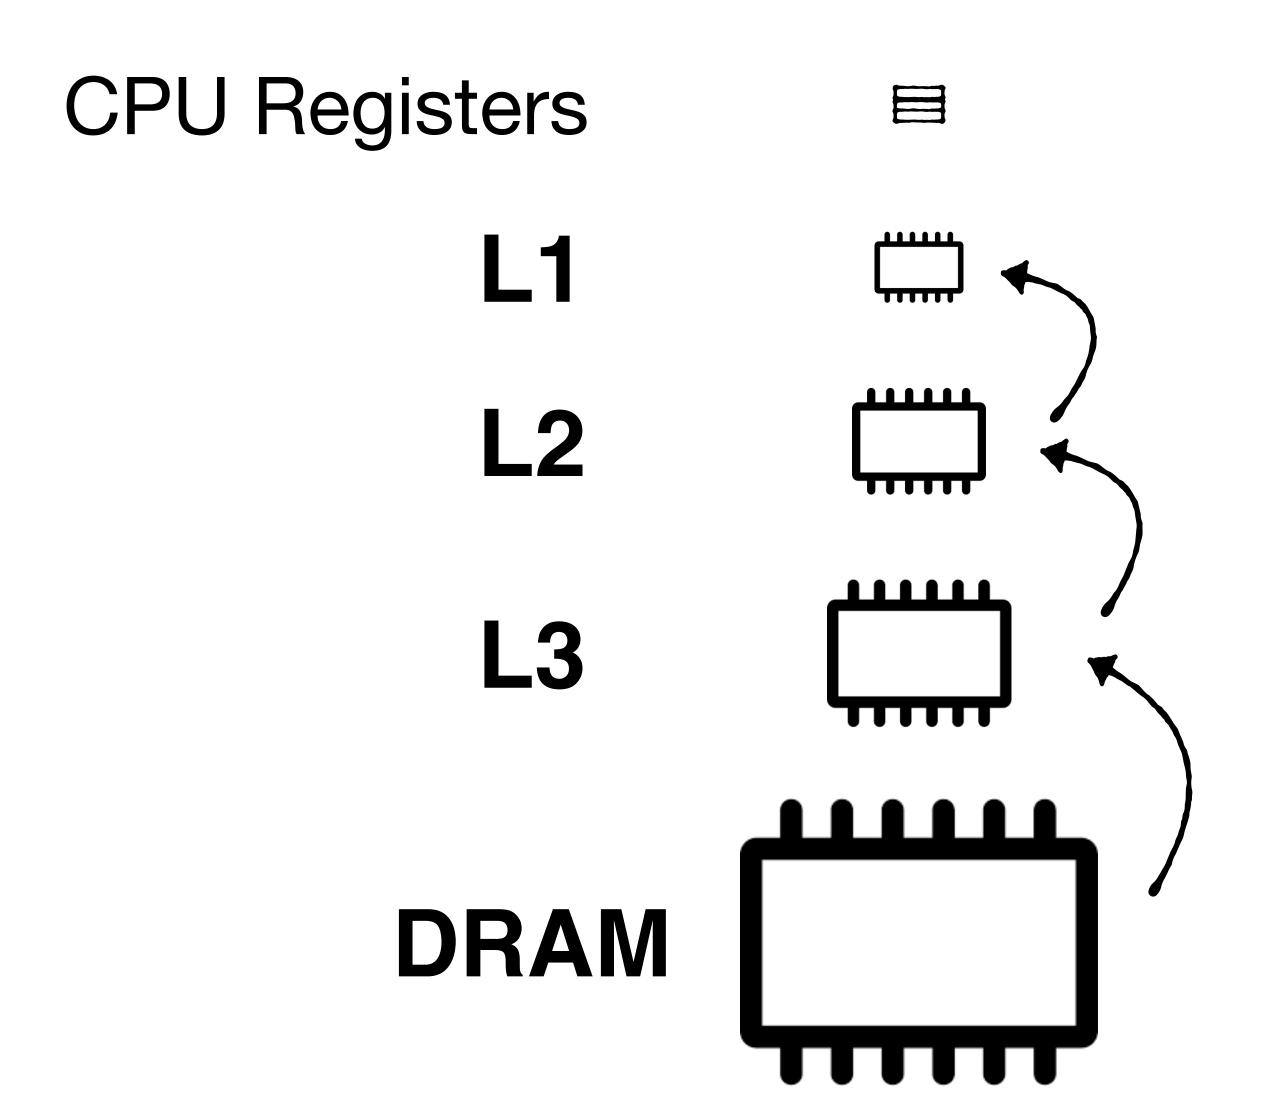


Avg. Negative Query = 2

Is this ok for modern hardware?

#### Recall the memory hierarchy

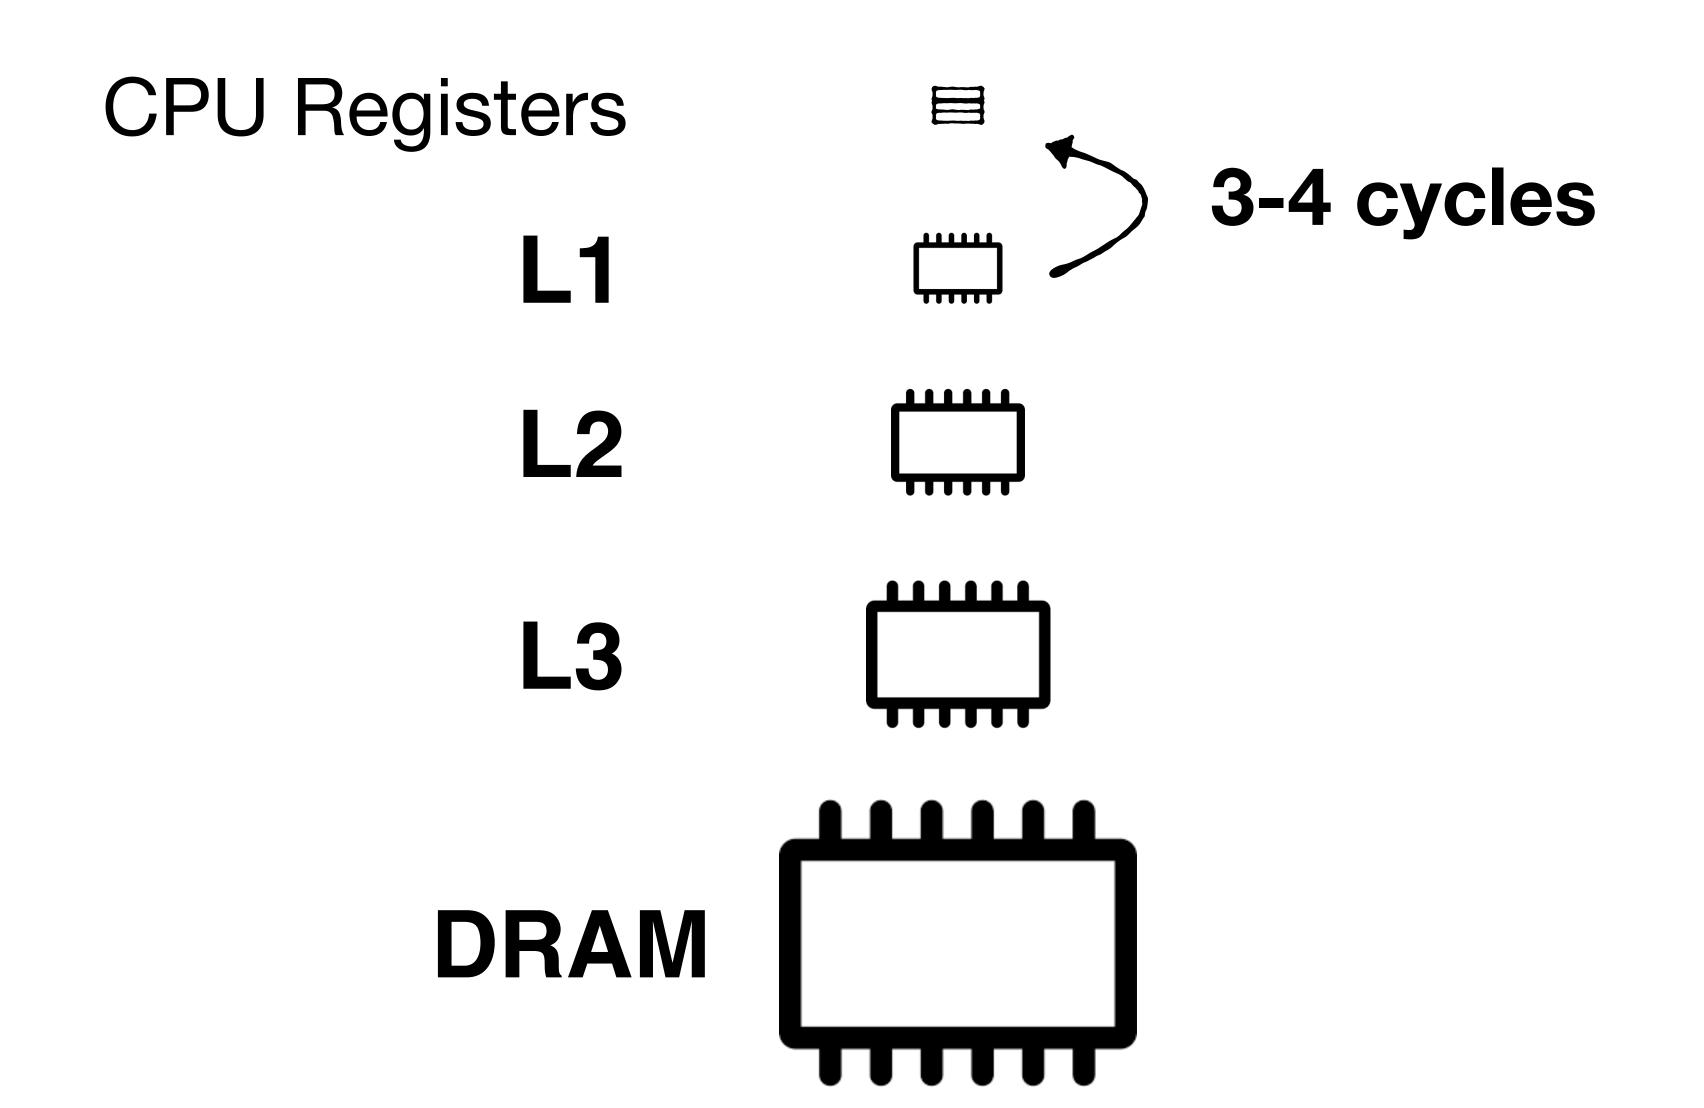


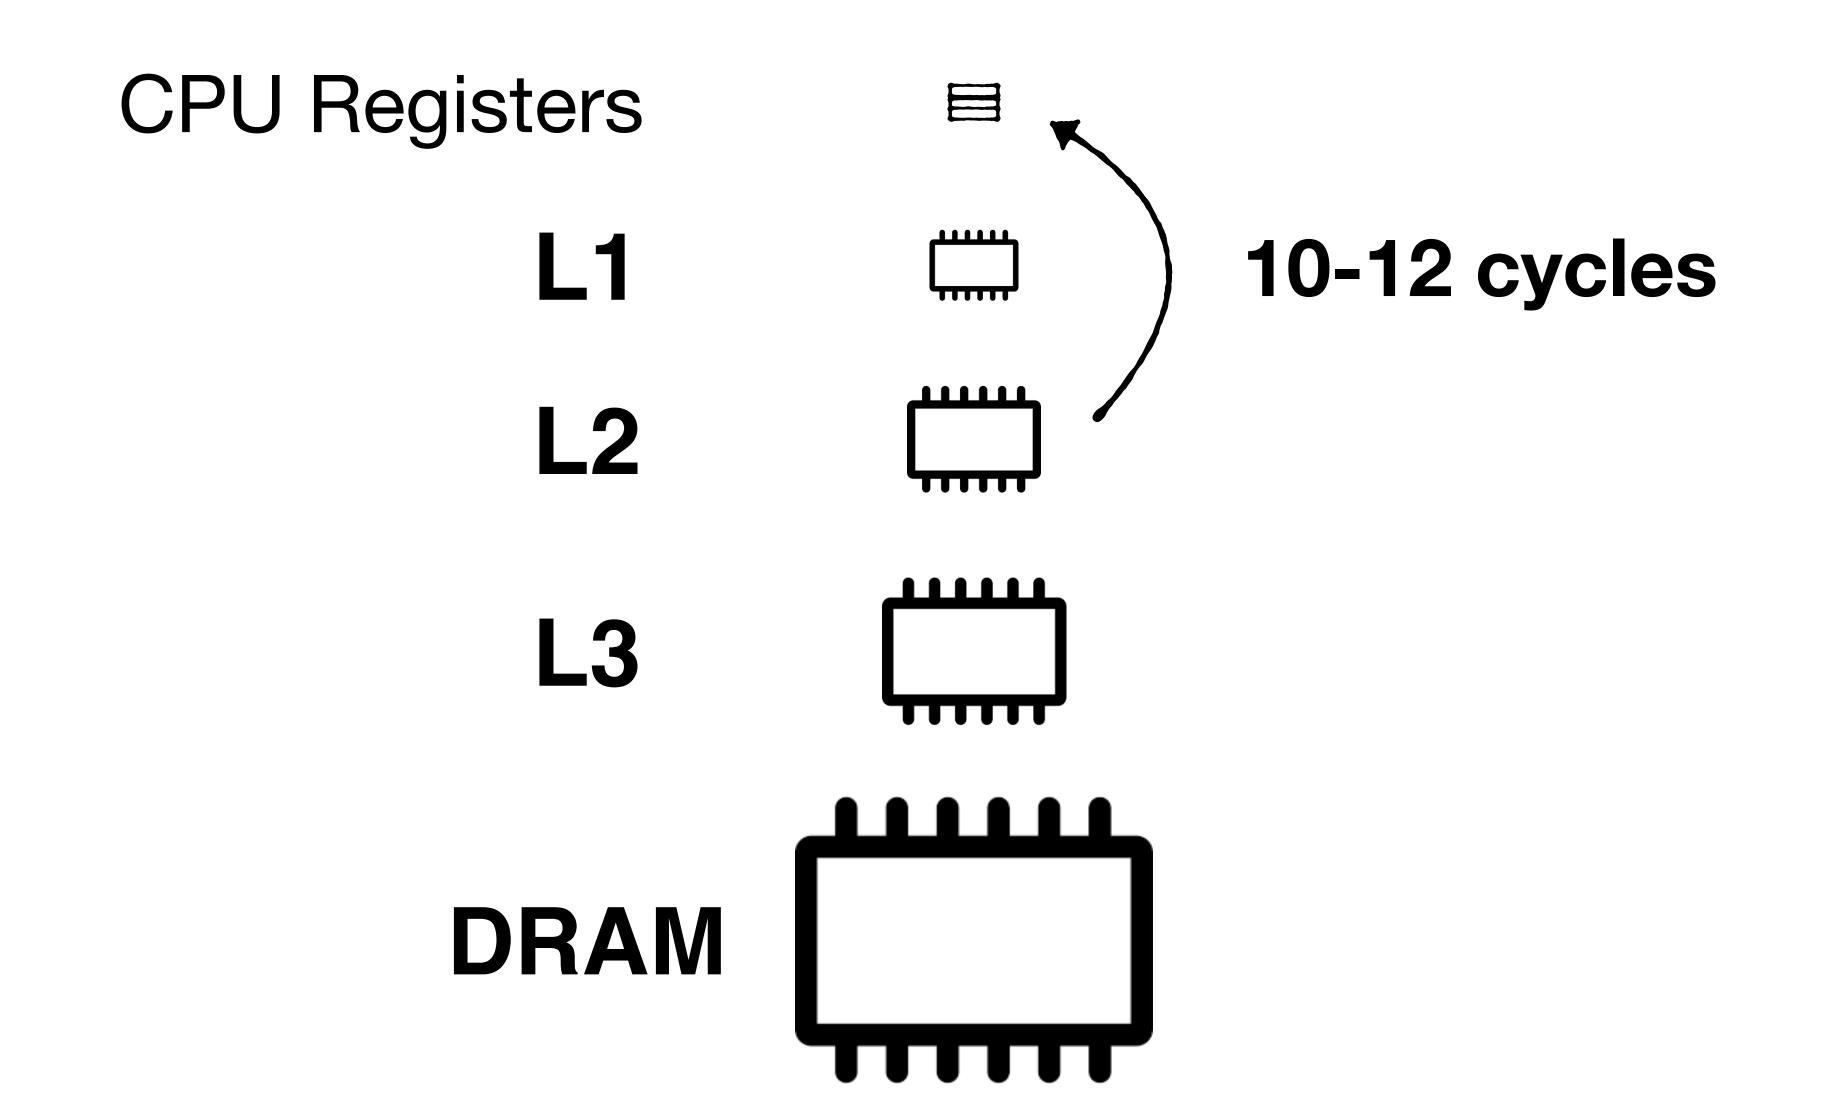


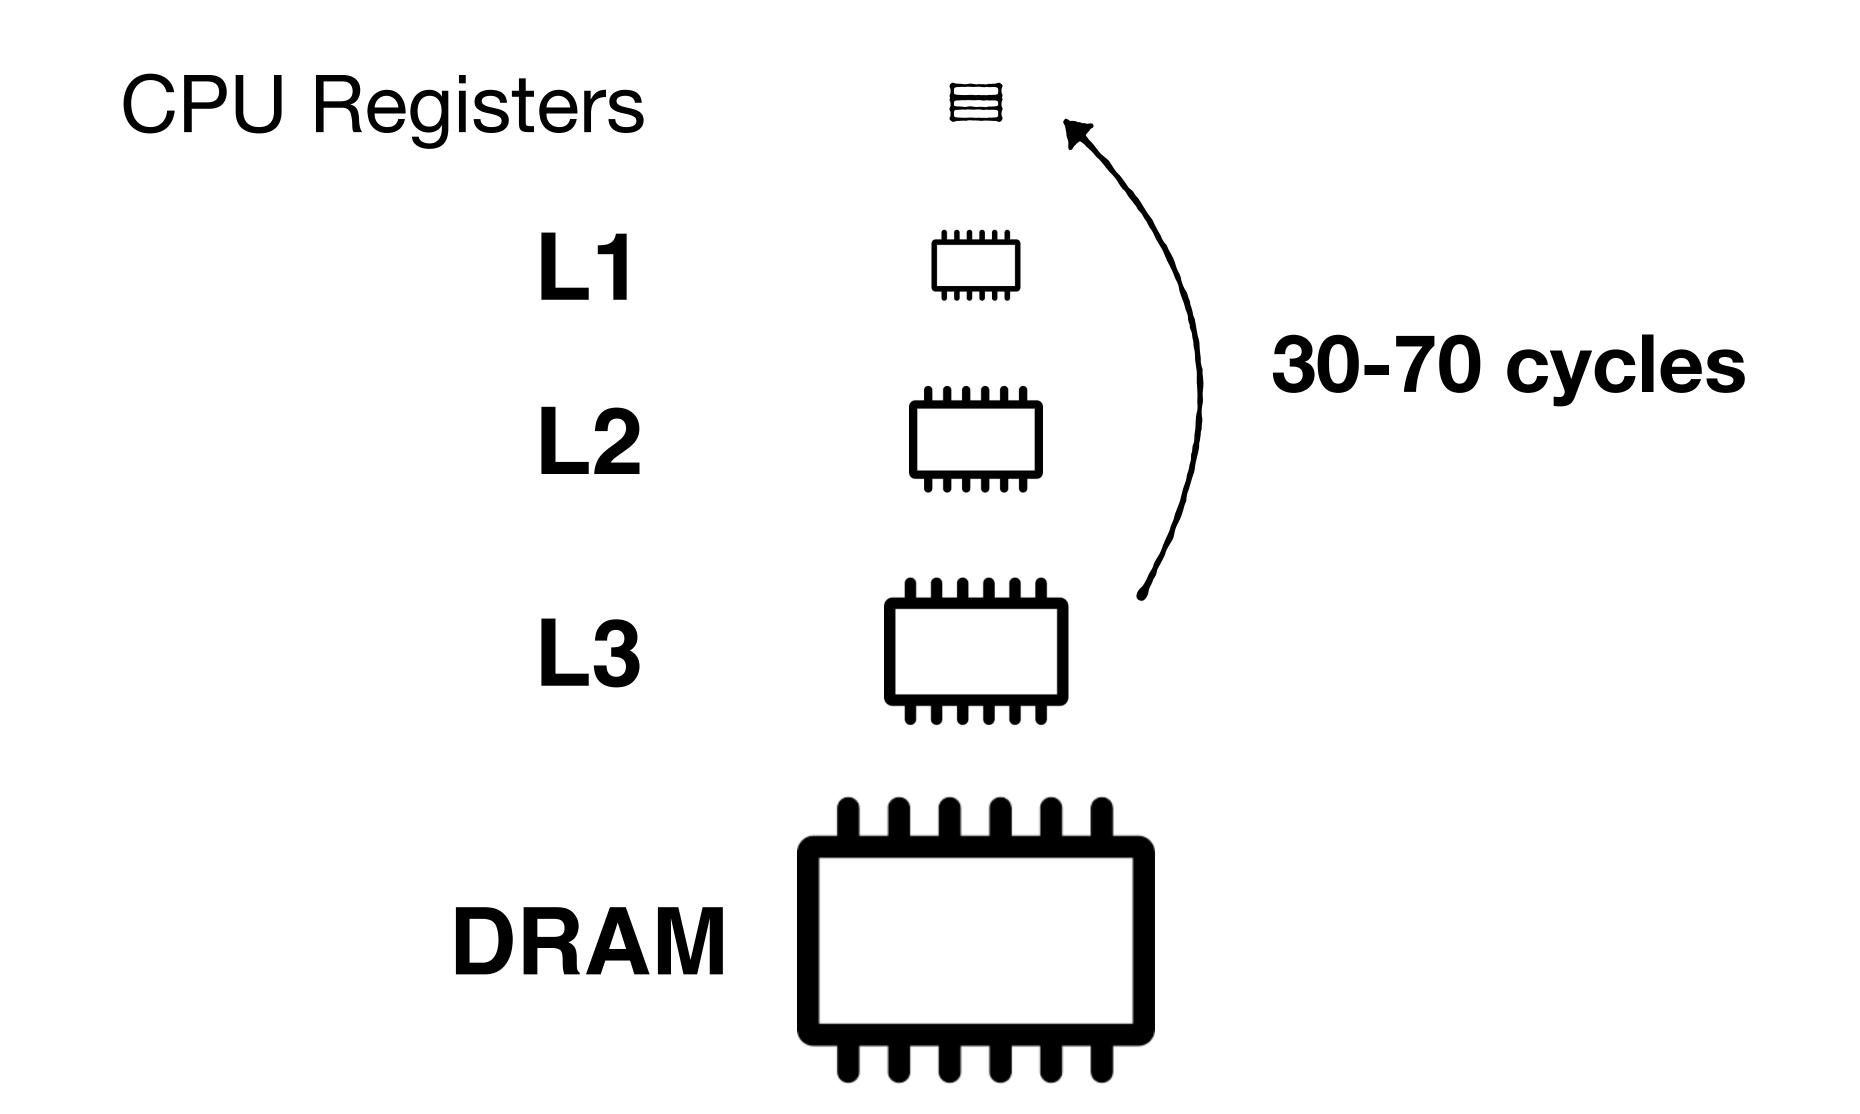
Move data at "cache line" granularity (e.g., 64B)

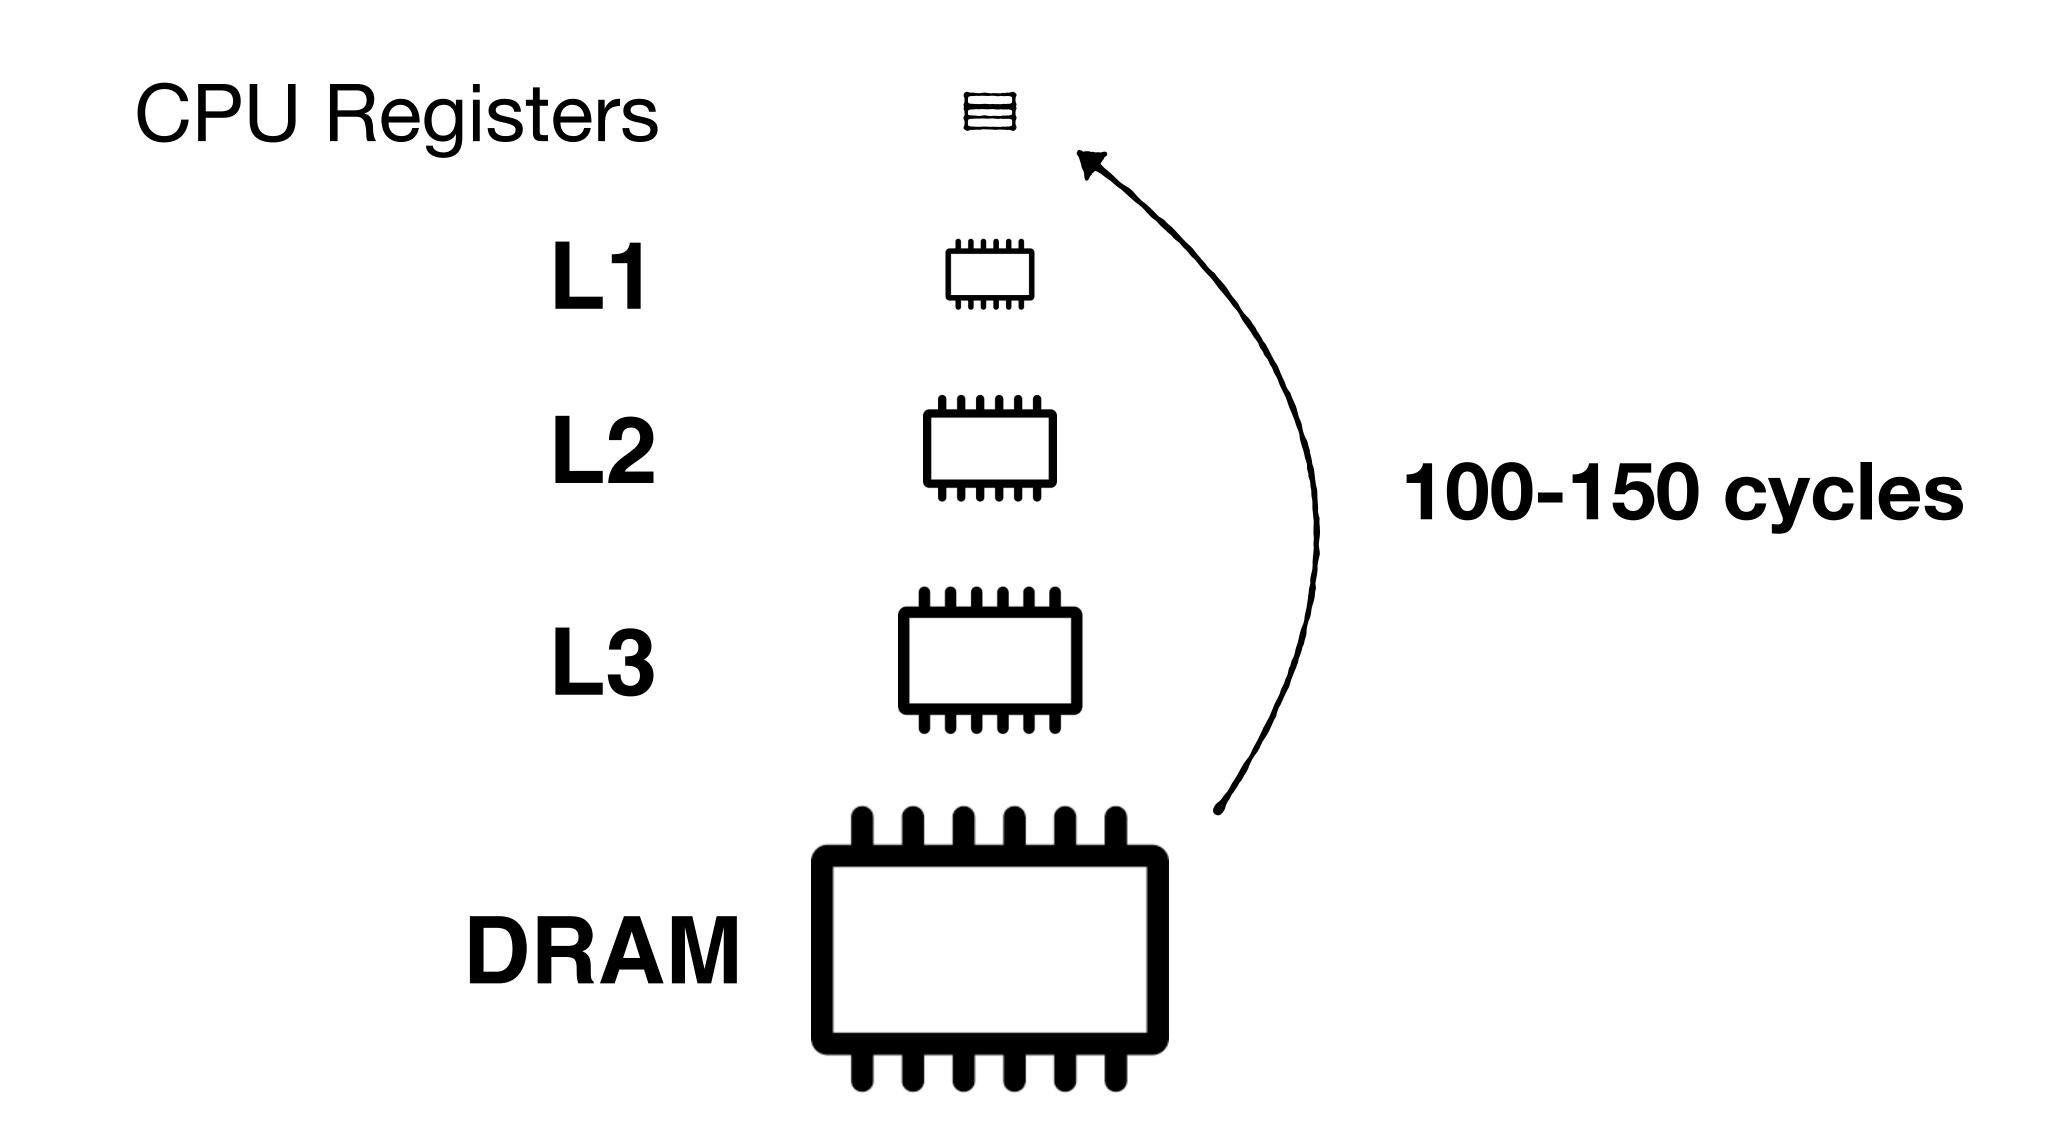
# **CPU Registers** DRAM

## Move data at "word" granularity (e.g., 8B)

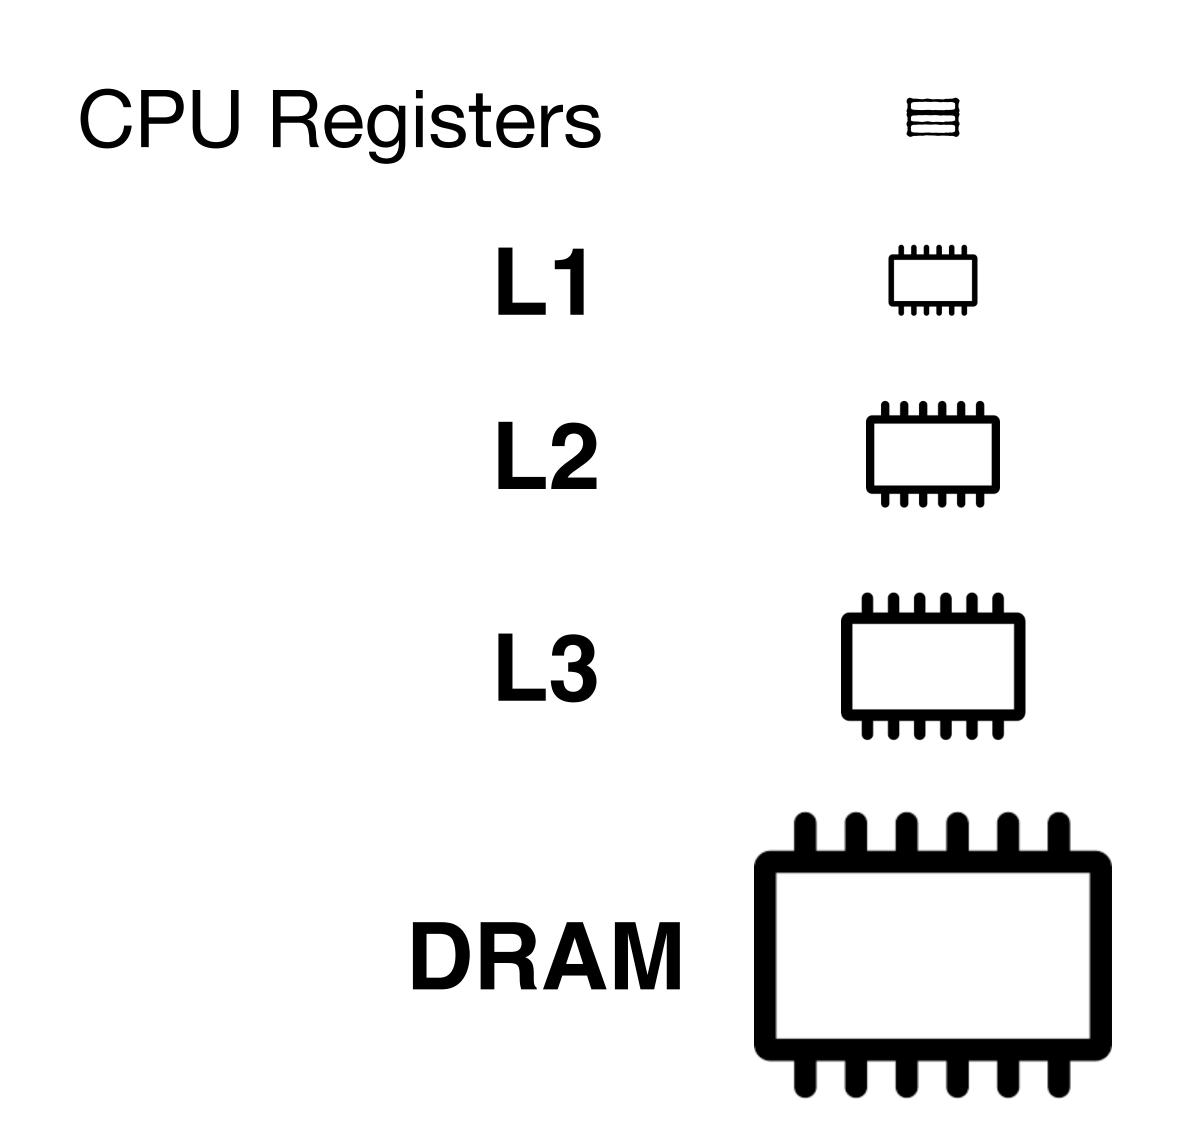








**Source:** http://ithare.com/infographics-operation-costs-in-cpu-clock-cycles/ (Numbers from **2016**)



#### Each hash function can lead to a cache miss

Insertion =  $M/N \cdot ln(2)$ 

Positive Query =  $M/N \cdot ln(2)$ 

Avg. Negative Query = 2

#### Each hash function can lead to a cache miss

Insertion =  $M/N \cdot ln(2) \cdot 100 ns$ 

Positive Query =  $M/N \cdot ln(2) \cdot 100 ns$ 

Avg. Negative Query = 2 · 100 ns

#### Observation

### Basic Bloom filter is slowest filter



#### Observation

### Basic Bloom filter is slowest filter



## With hardware optimizations, it is the fastest



#### Observation

Basic Bloom filter is slowest filter

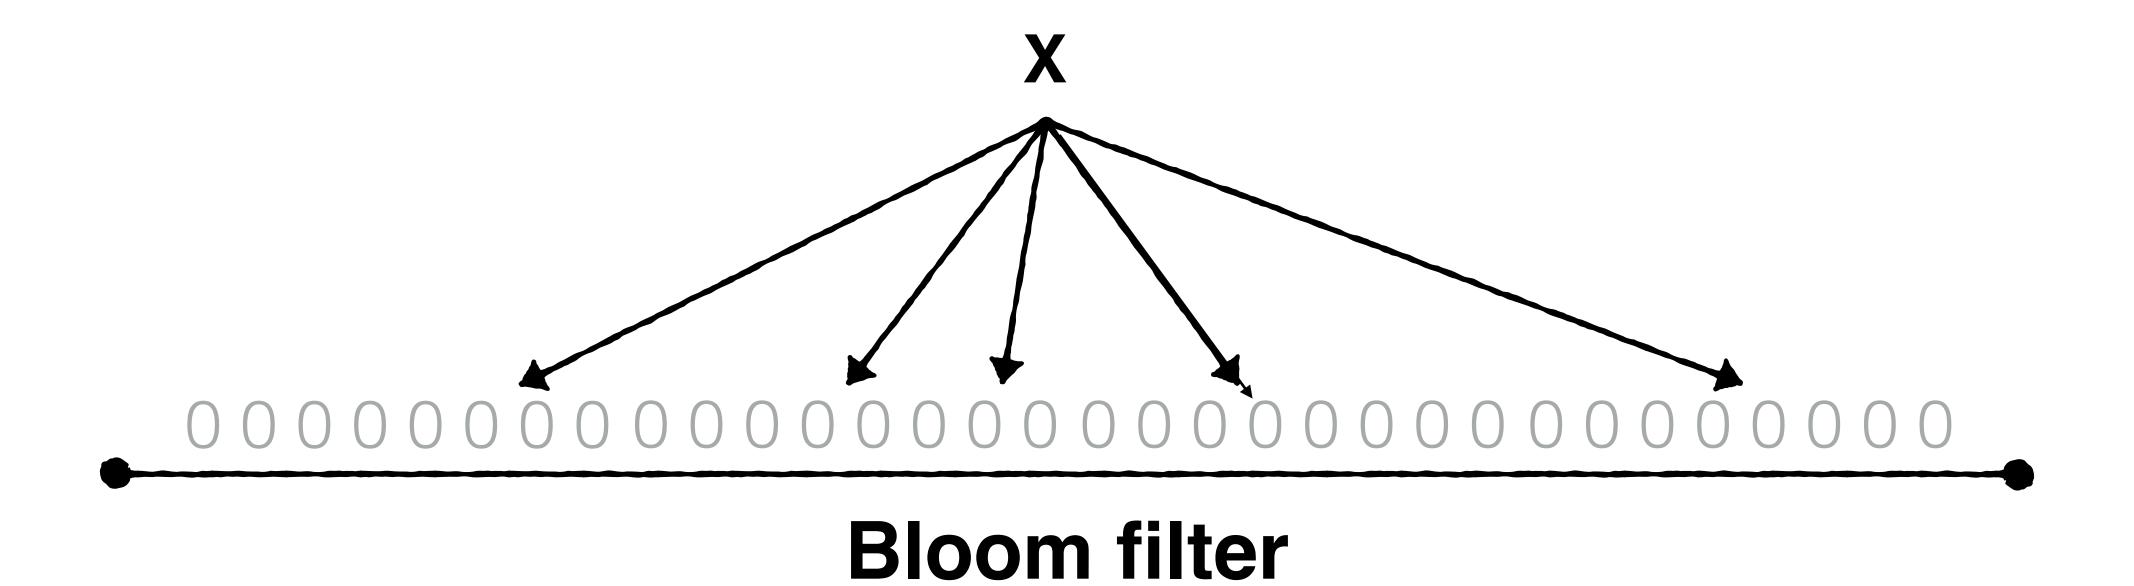
With hardware optimizations, it is the fastest



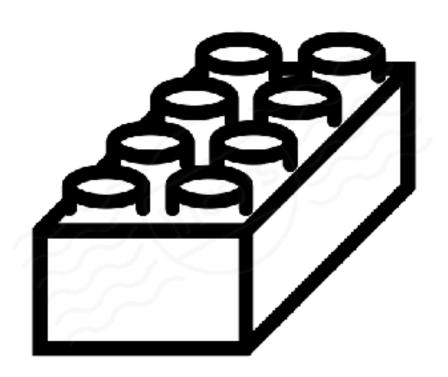


lends itself to hardware optimization

#### Alleviate random accesses?



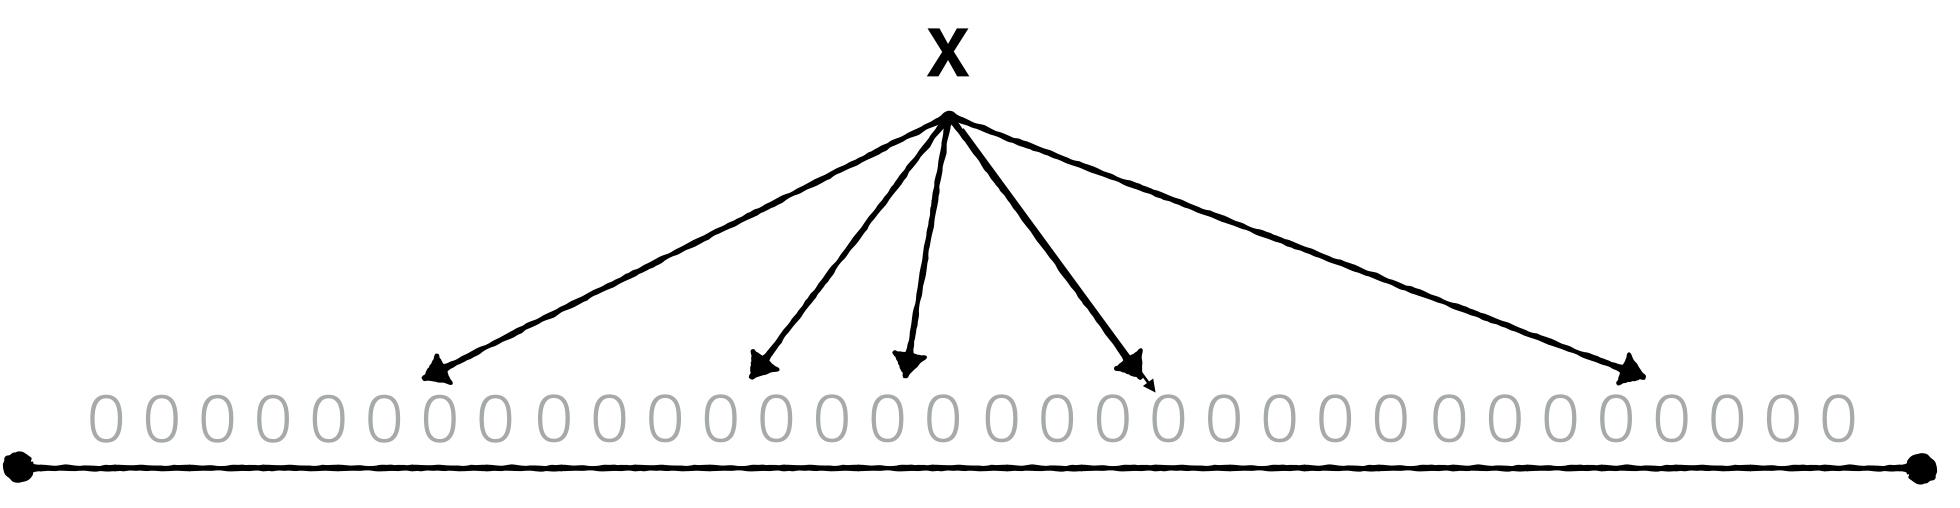
#### **Blocked Bloom Filters**



#### Cache-, Hash- and Space-Efficient Bloom Filters

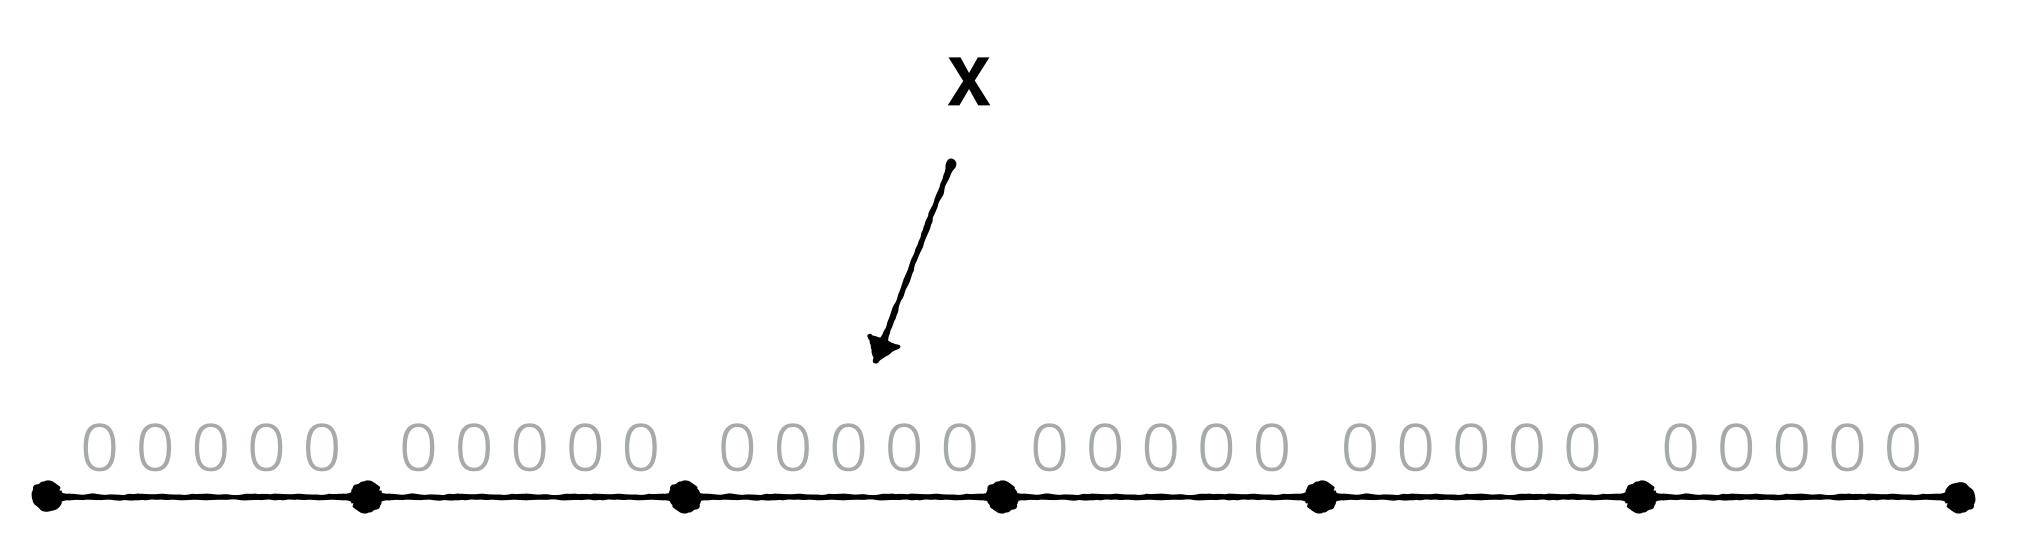
Journal of Experimental Algorithms, 2010

Felix Putze, Peter Sanders, Johannes Singler

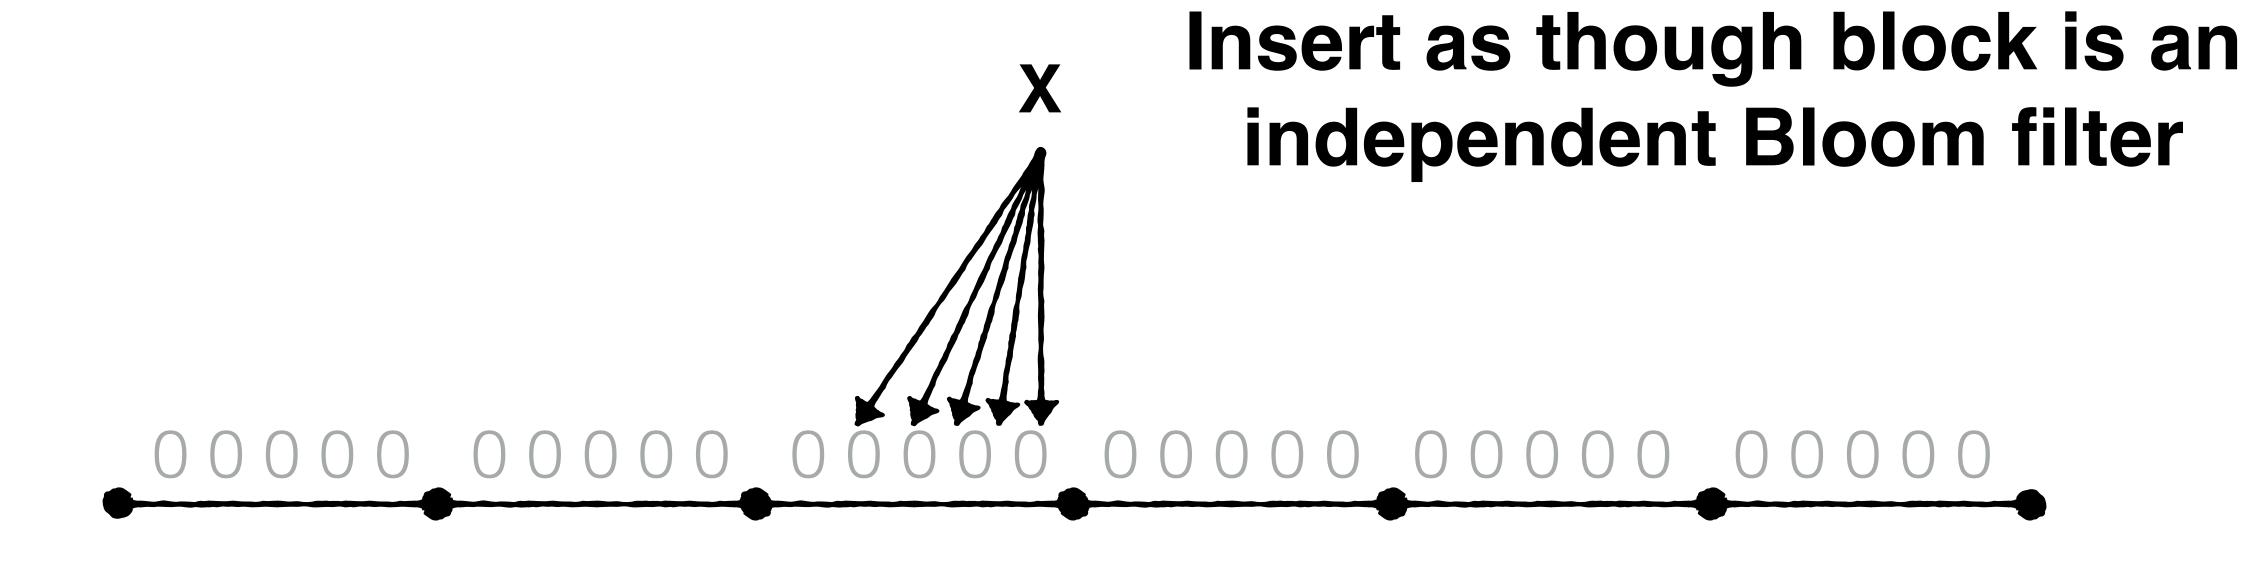


**Bloom filter** 

#### Hash to one block, sized as a cache line

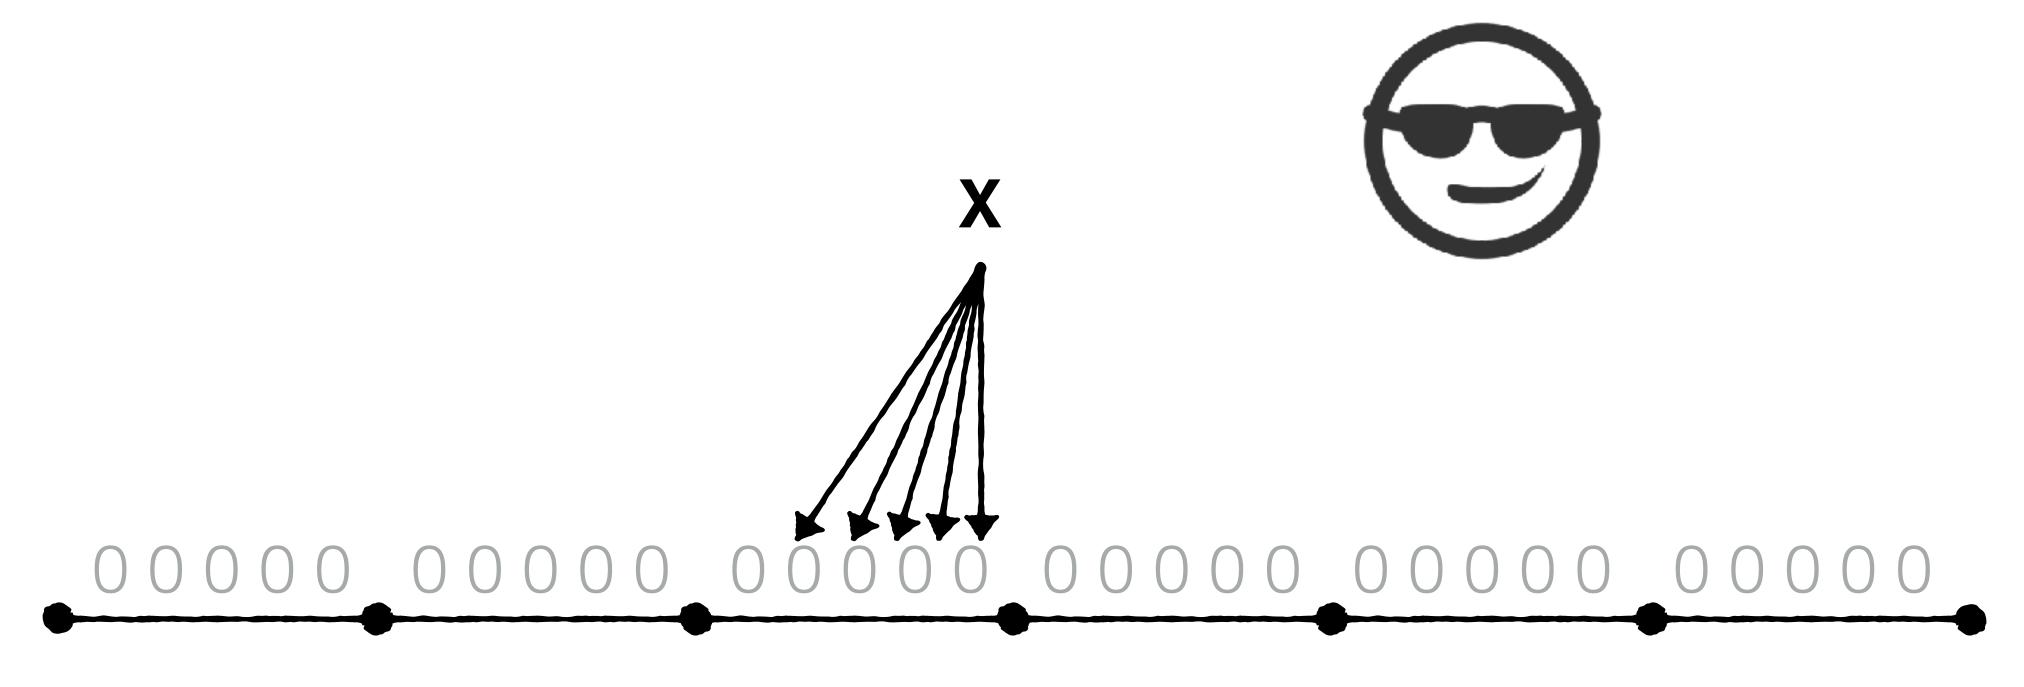


Blocked Bloom filter



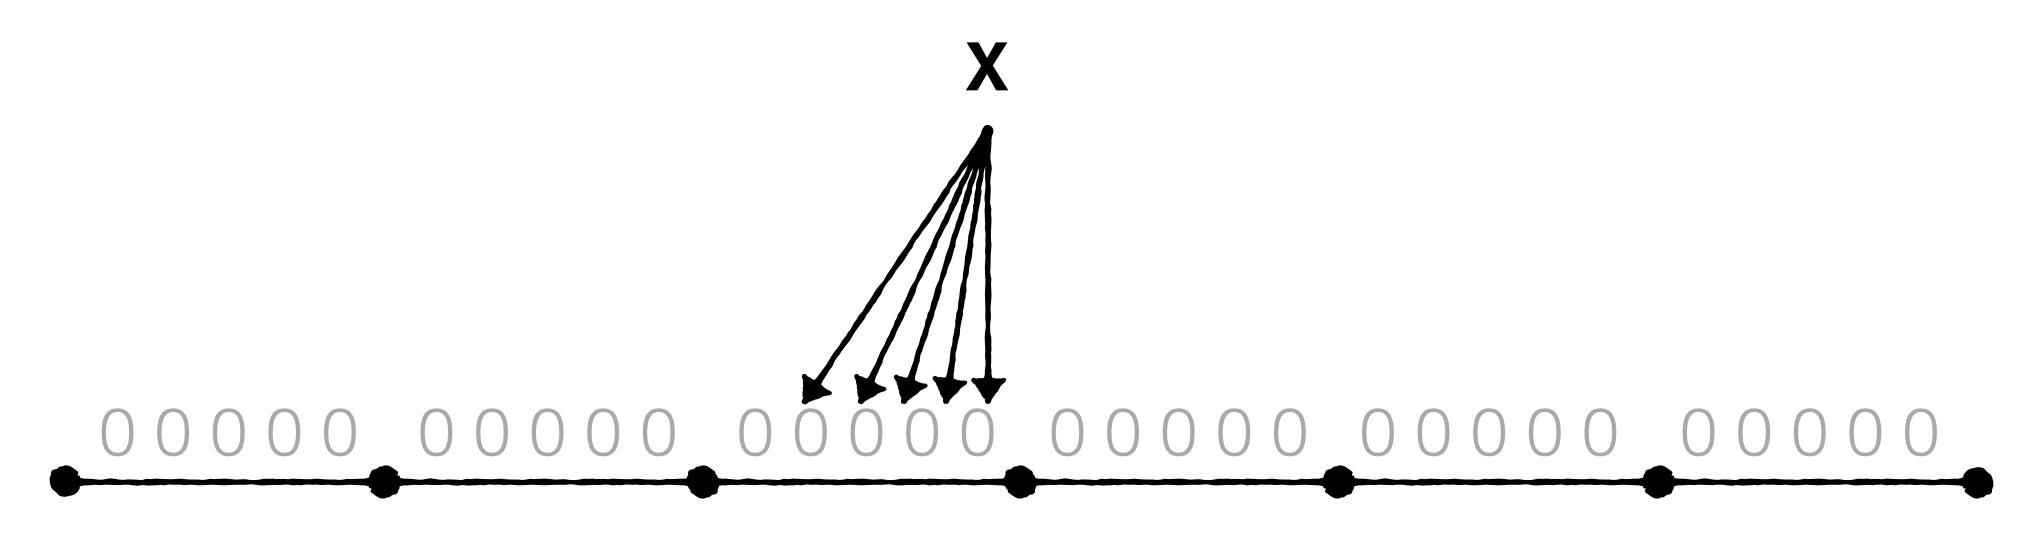
Blocked Bloom filter

#### 1 cache miss per query/insertion



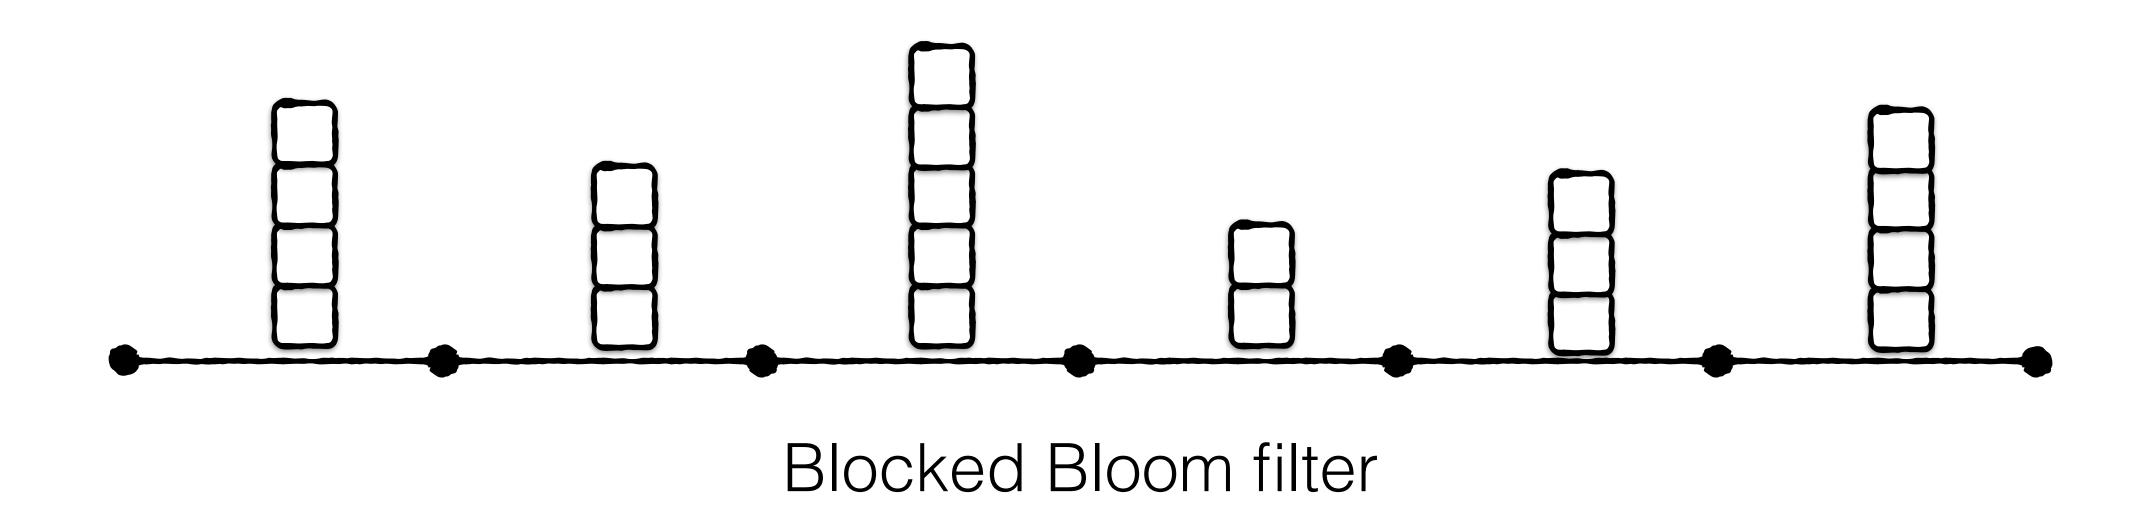
Blocked Bloom filter

### 1 cache miss per query/insertion Anything bad?

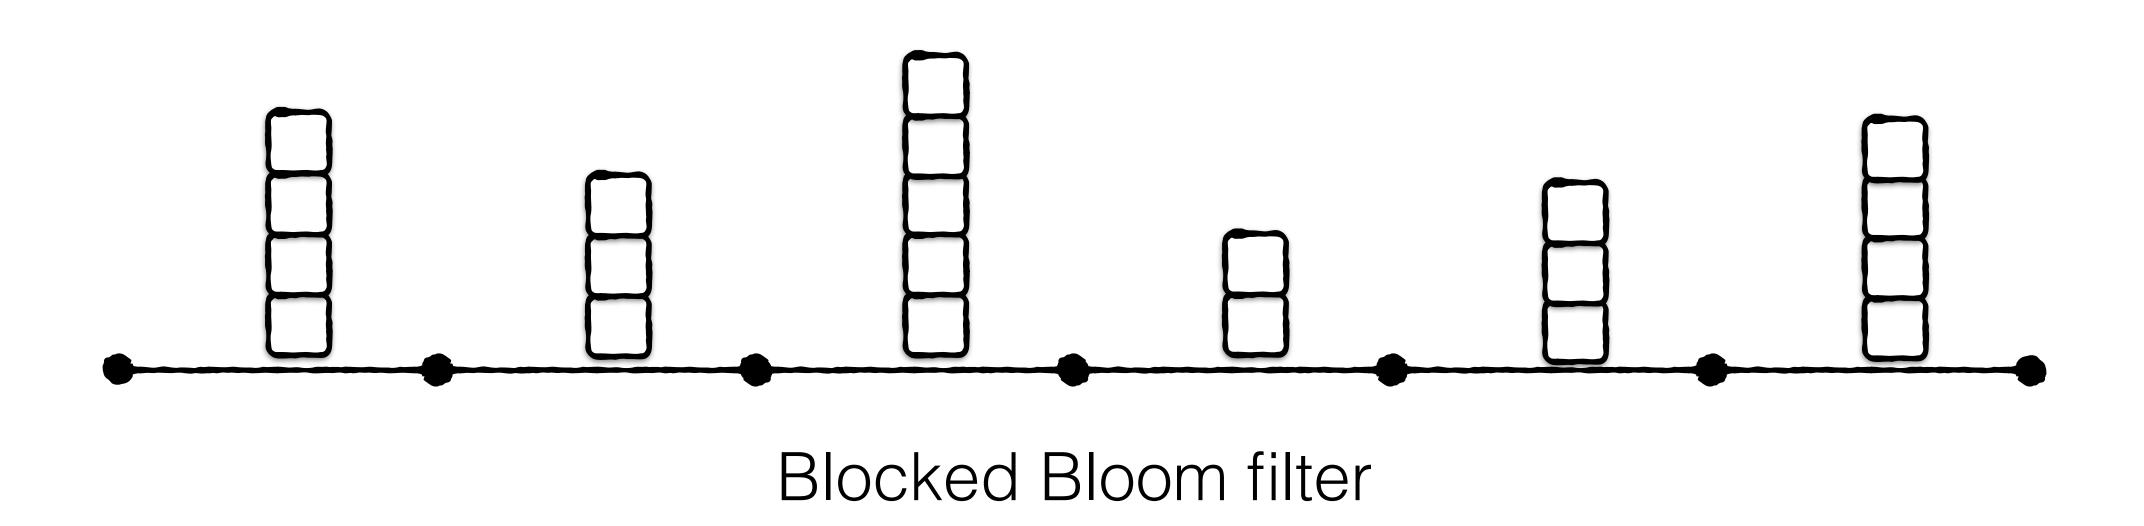


Blocked Bloom filter

#### uneven distribution of entries across blocks

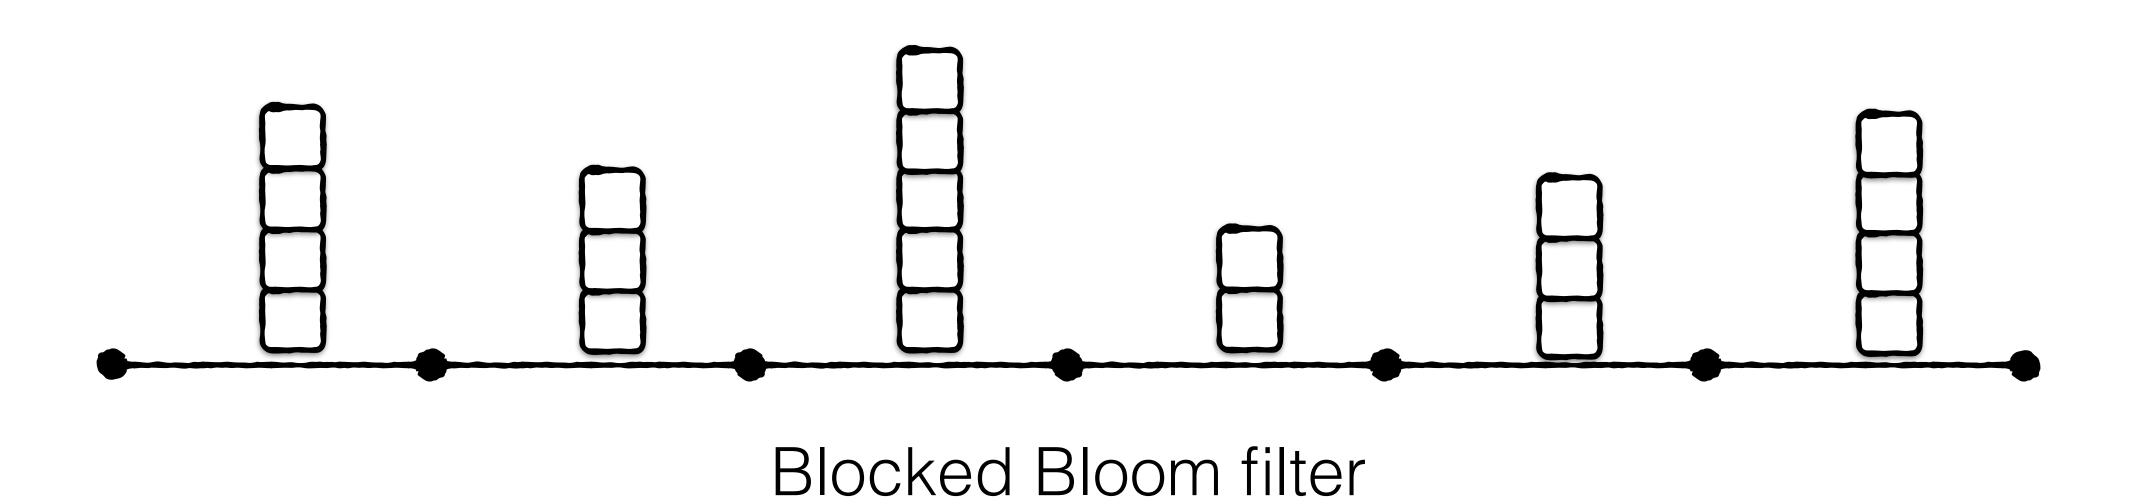


### uneven distribution of entries across blocks impact on FPR?

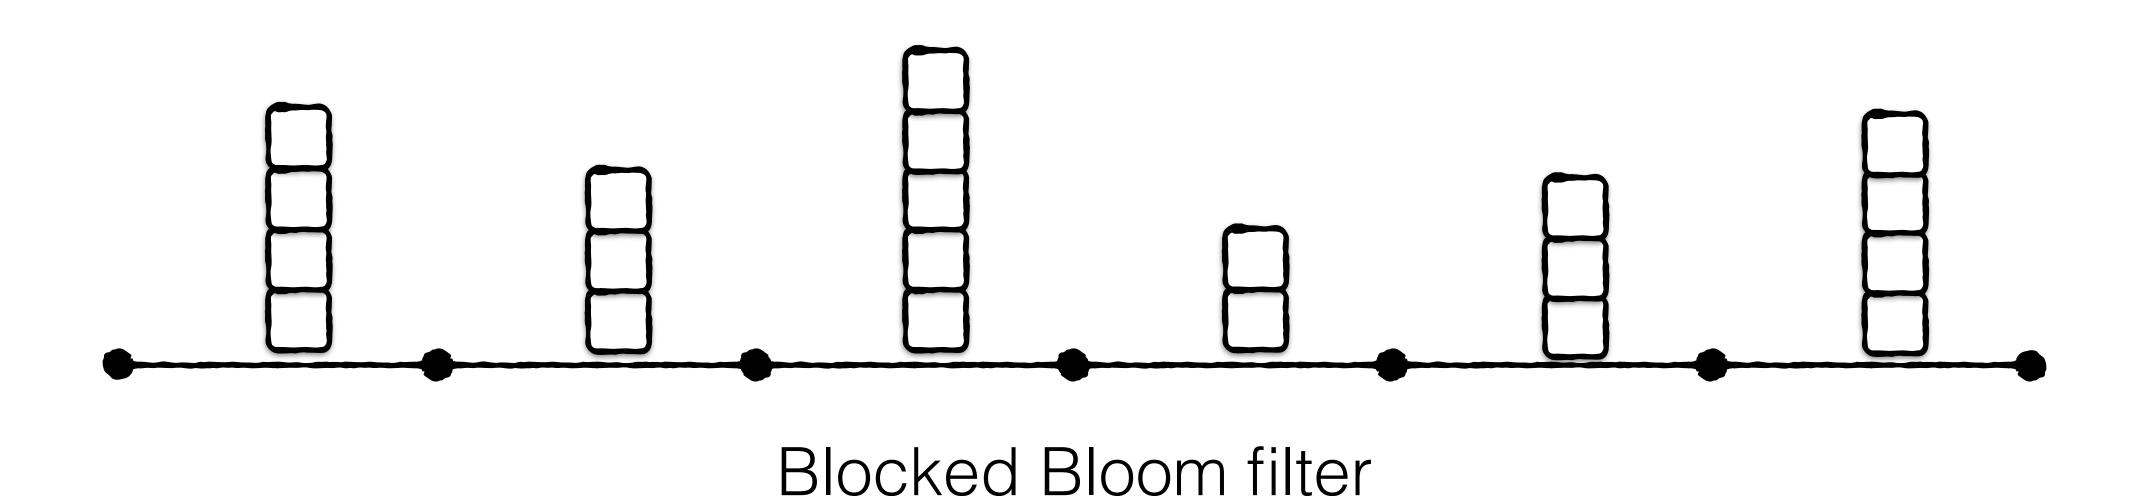


uneven distribution of entries across blocks impact on FPR?

#### General analysis technique for hash tables

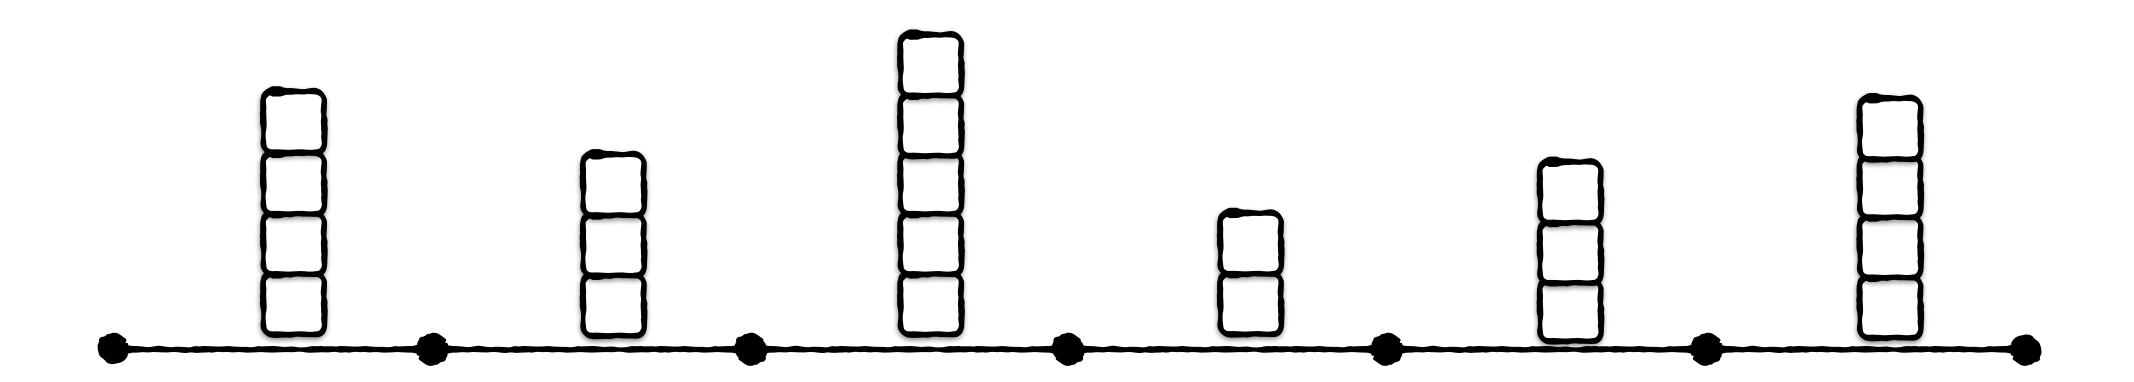


$$f(n, p, i) = \binom{n}{i} \cdot p^{i} \cdot (1 - p)^{n - i}$$



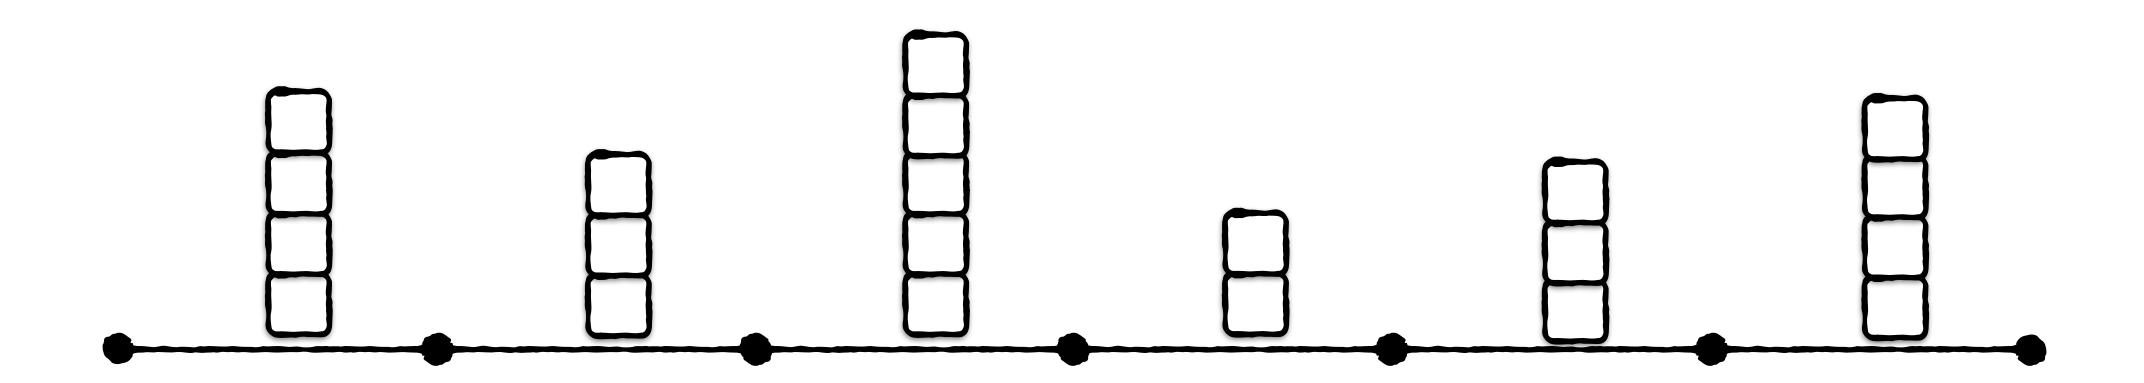
$$f(n, p, i) = \binom{n}{i} \cdot p^{i} \cdot (1 - p)^{n - i}$$

# entries



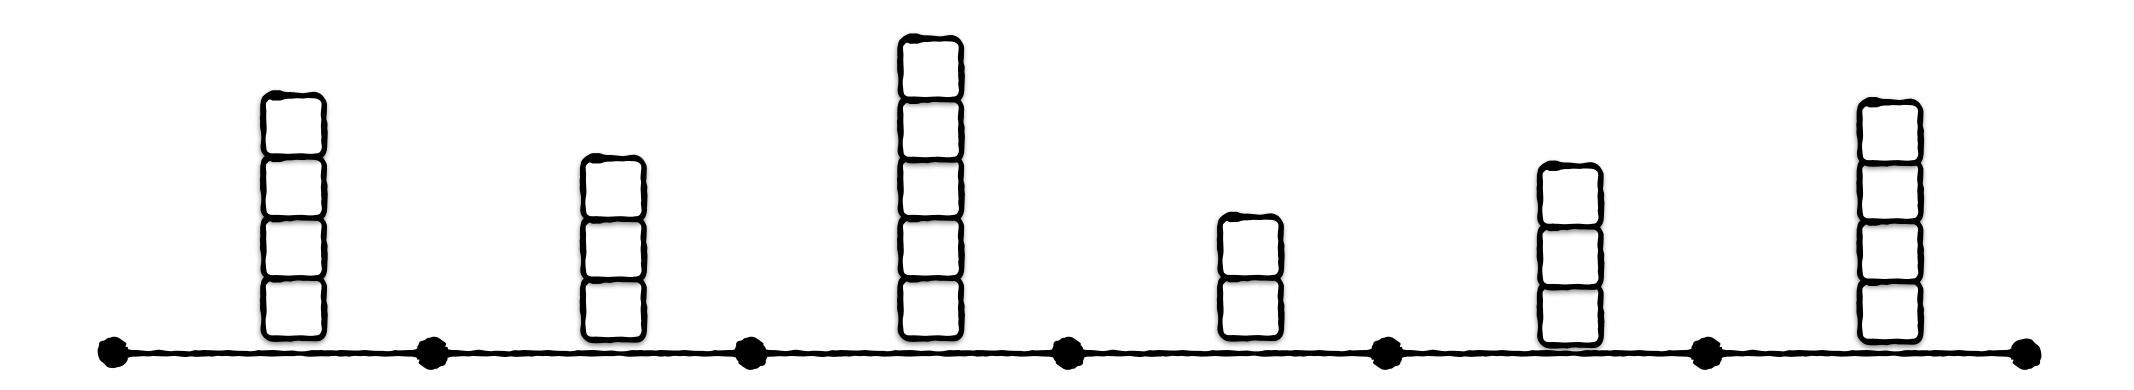
$$f(n, p, i) = \binom{n}{i} \cdot p^{i} \cdot (1 - p)^{n-i}$$

### Prob of 1 entry falling into a given block, i.e., 1/n



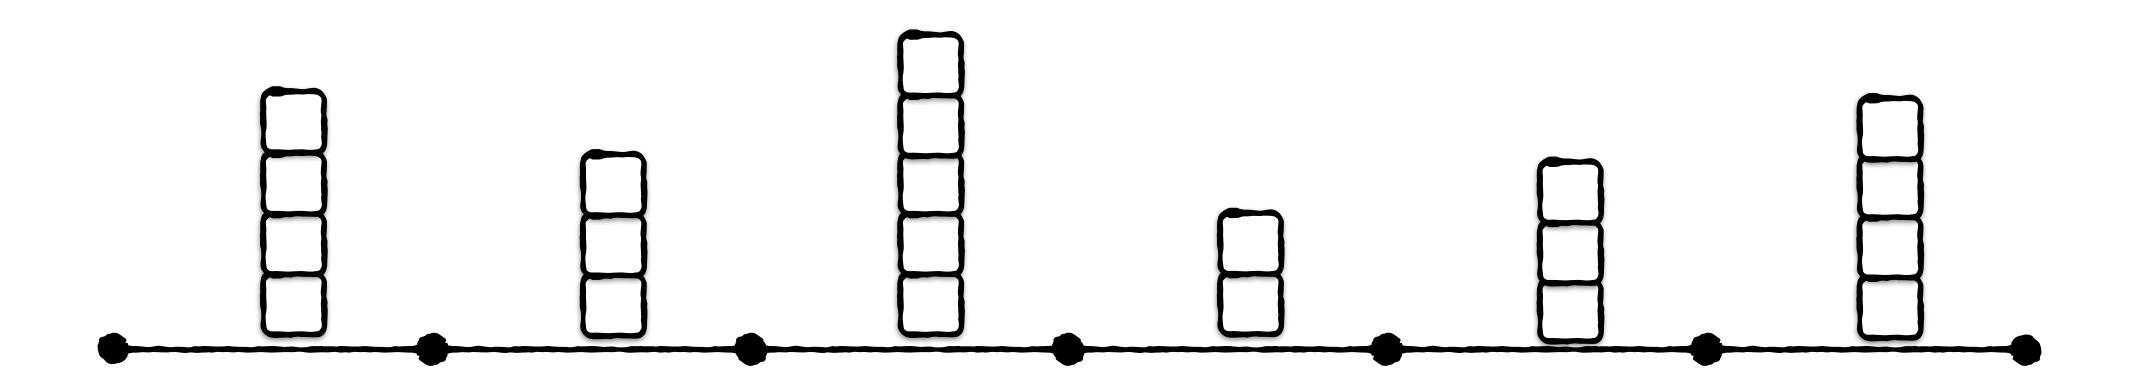
$$f(n, p, i) = \binom{n}{i} \cdot p^{i} \cdot (1 - p)^{n - i}$$

#### i entries falling into our bucket



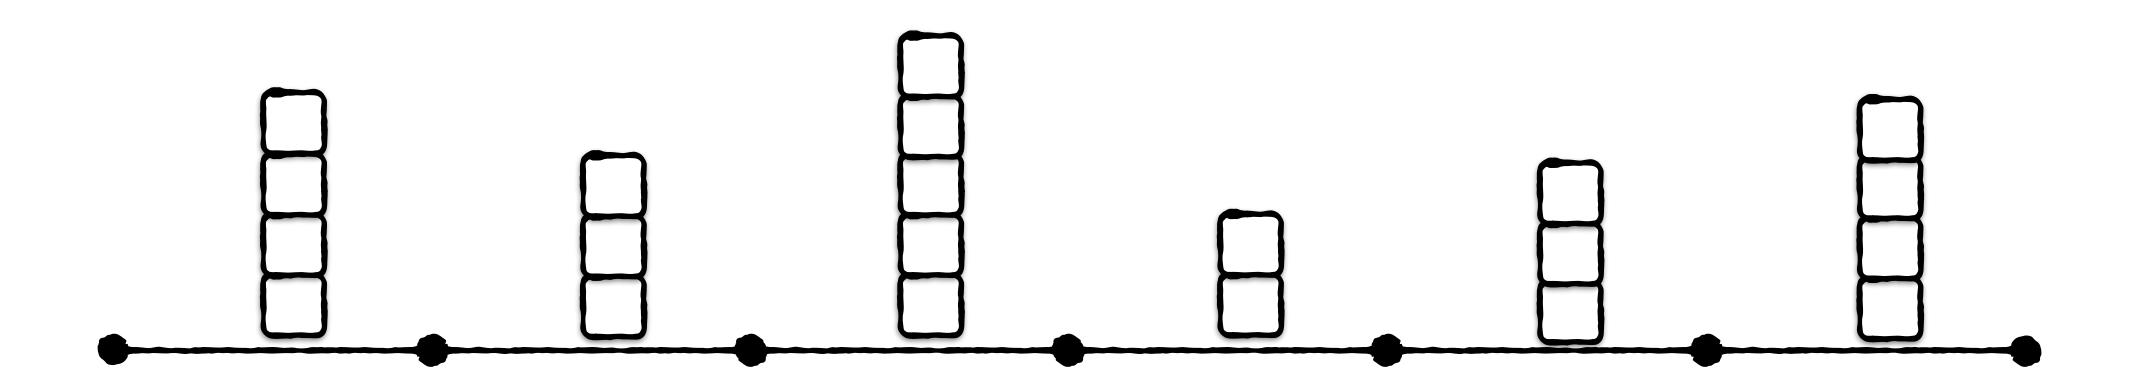
$$f(n, p, i) = \binom{n}{i} \cdot p^{i} \cdot (1 - p)^{n - i}$$

Ways of choosing k out of n entries



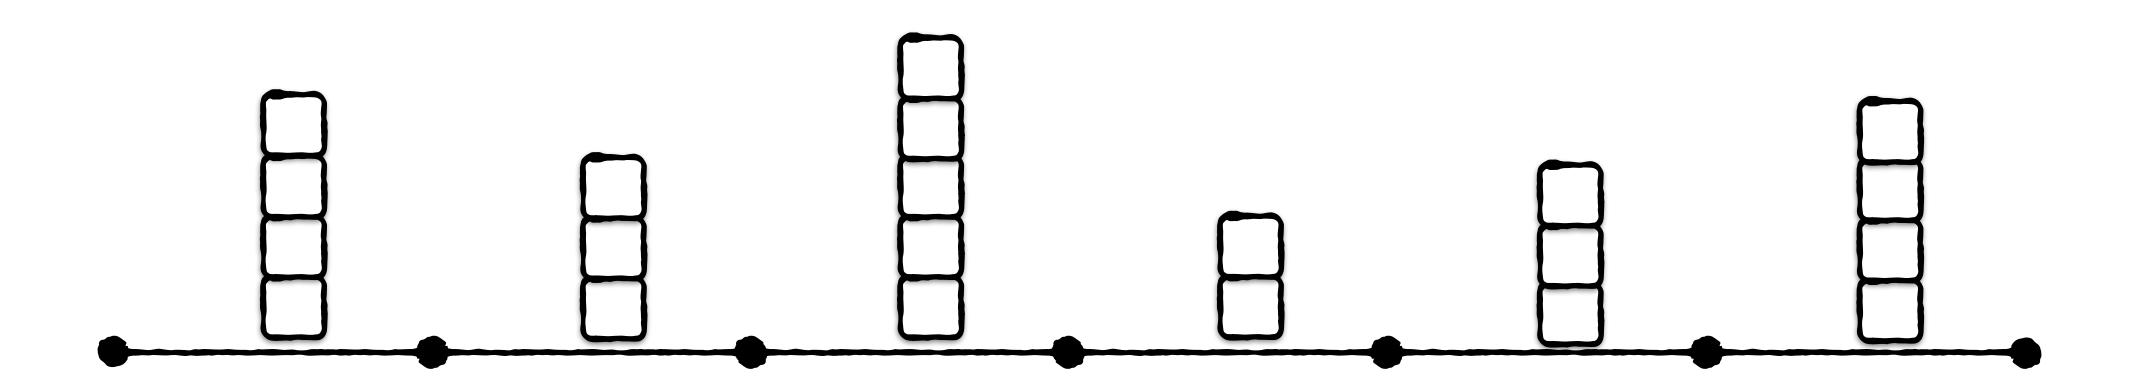
$$f(n, p, i) = \binom{n}{i} \cdot p^{i} \cdot (1 - p)^{n - i}$$

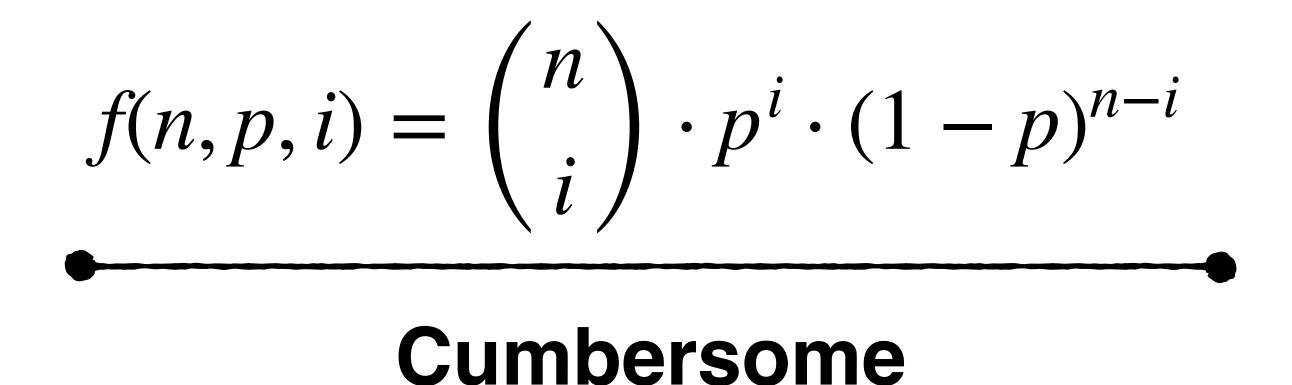
#### k entries falling into our block

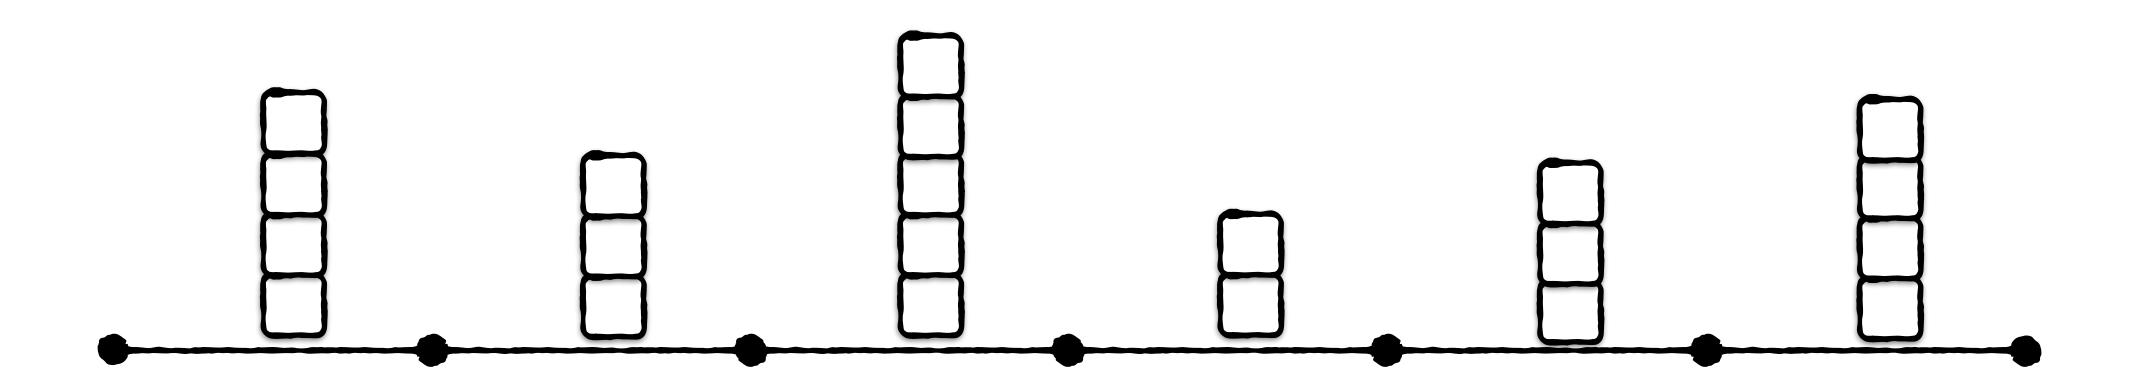


$$f(n, p, i) = \binom{n}{i} \cdot p^{i} \cdot (1 - p)^{n - i}$$

#### All other entries falling into other blocks



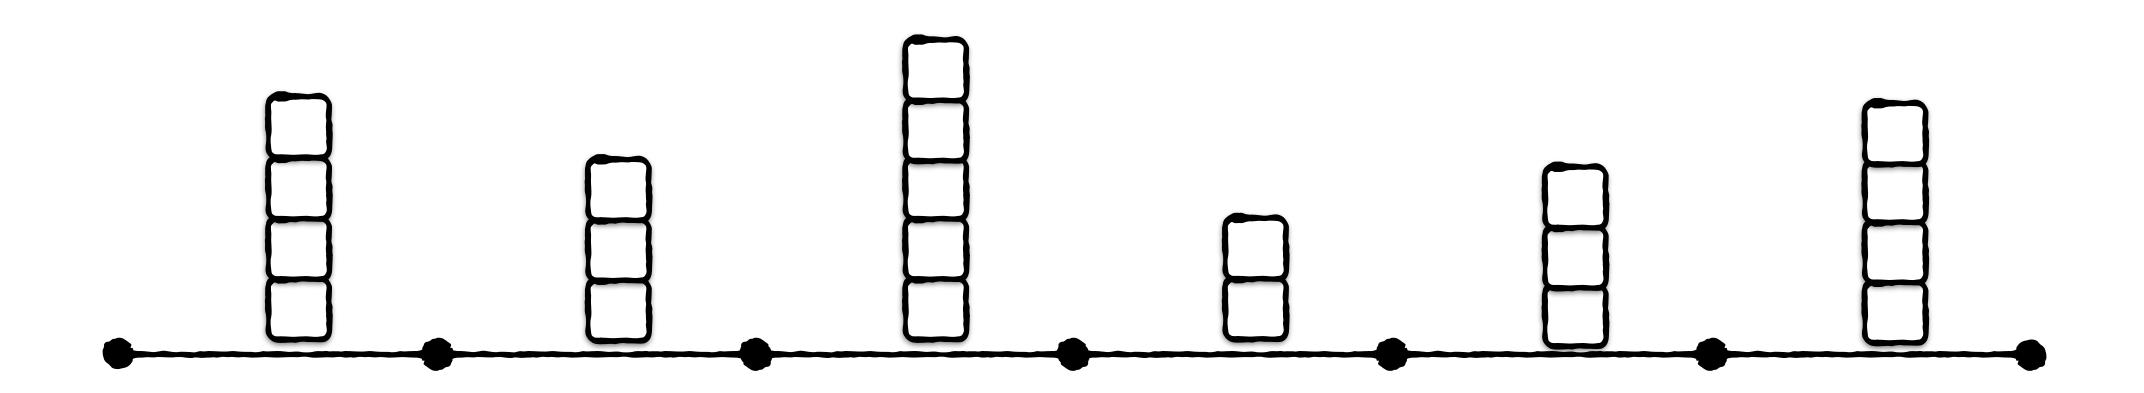




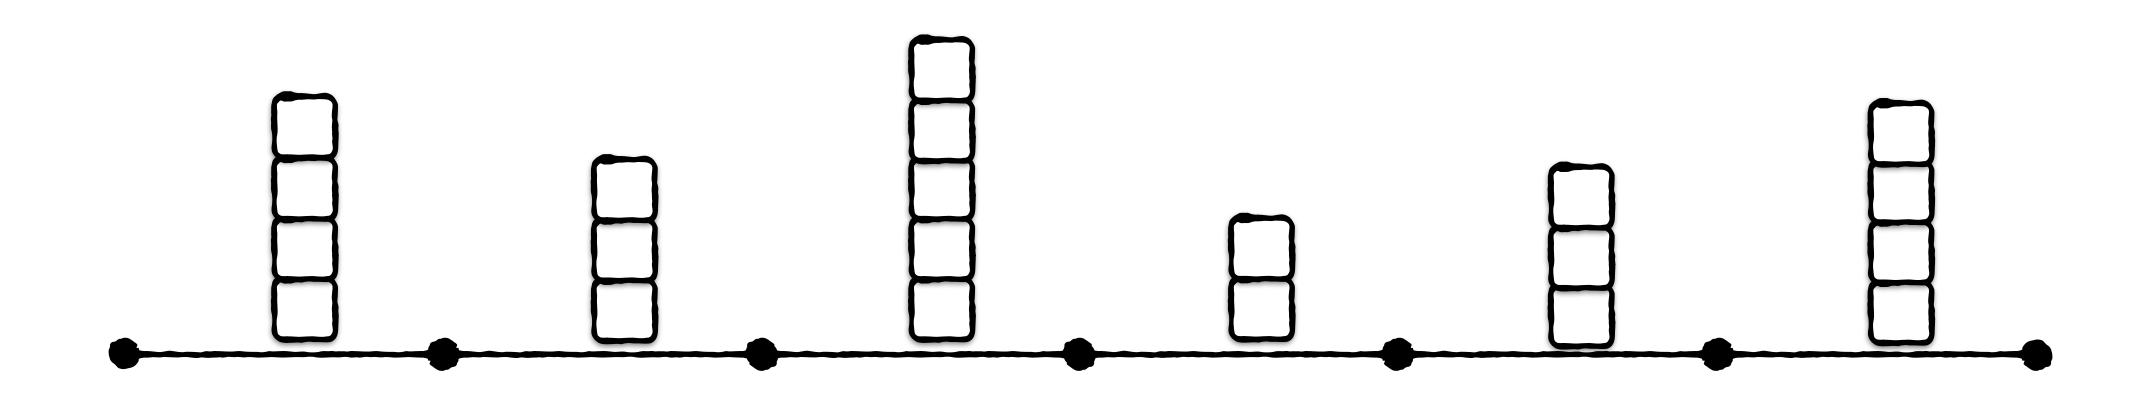
# entries per block follows binomial distribution

$$f(n, p, i) = \binom{n}{i} \cdot p^{i} \cdot (1 - p)^{n - i}$$

# For n → ∞, binomial converges to Poisson

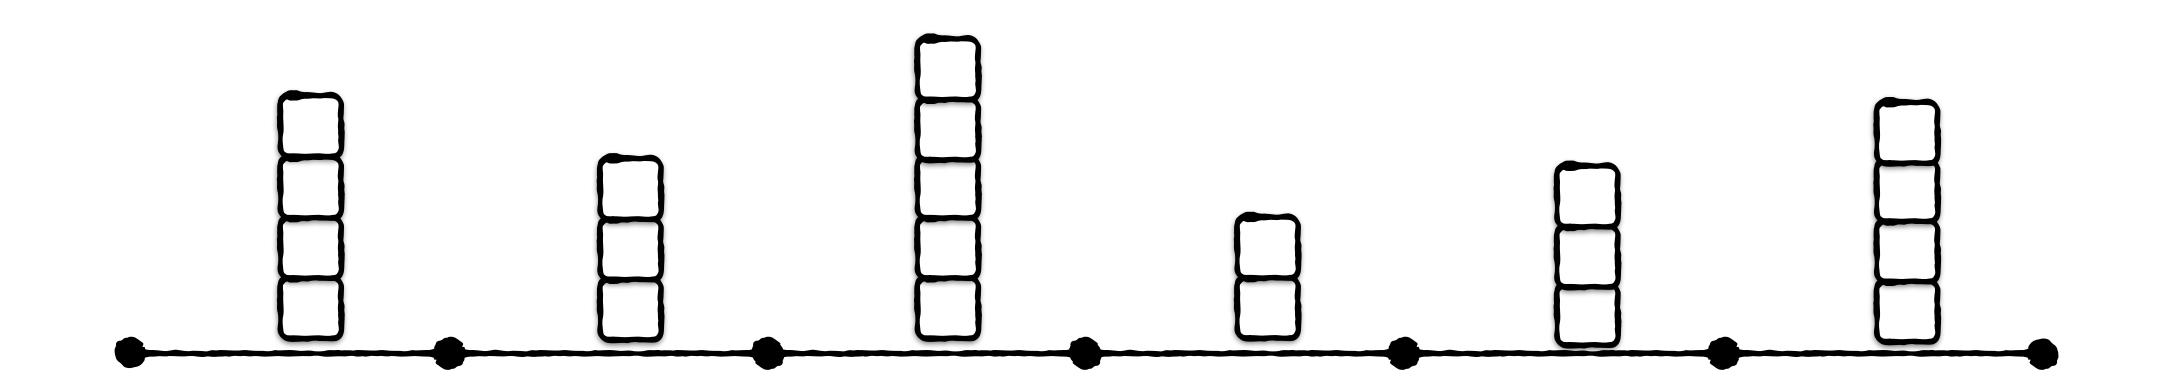


# P[i entries fall into given block] ~ Poisson(i, λ) = $\frac{\lambda^{i} \cdot e^{-\lambda}}{i!}$



P[i entries fall into given block] ~ Poisson(i,  $\lambda$ ) =  $\frac{\lambda^i \cdot e^{-\lambda}}{i!}$ 

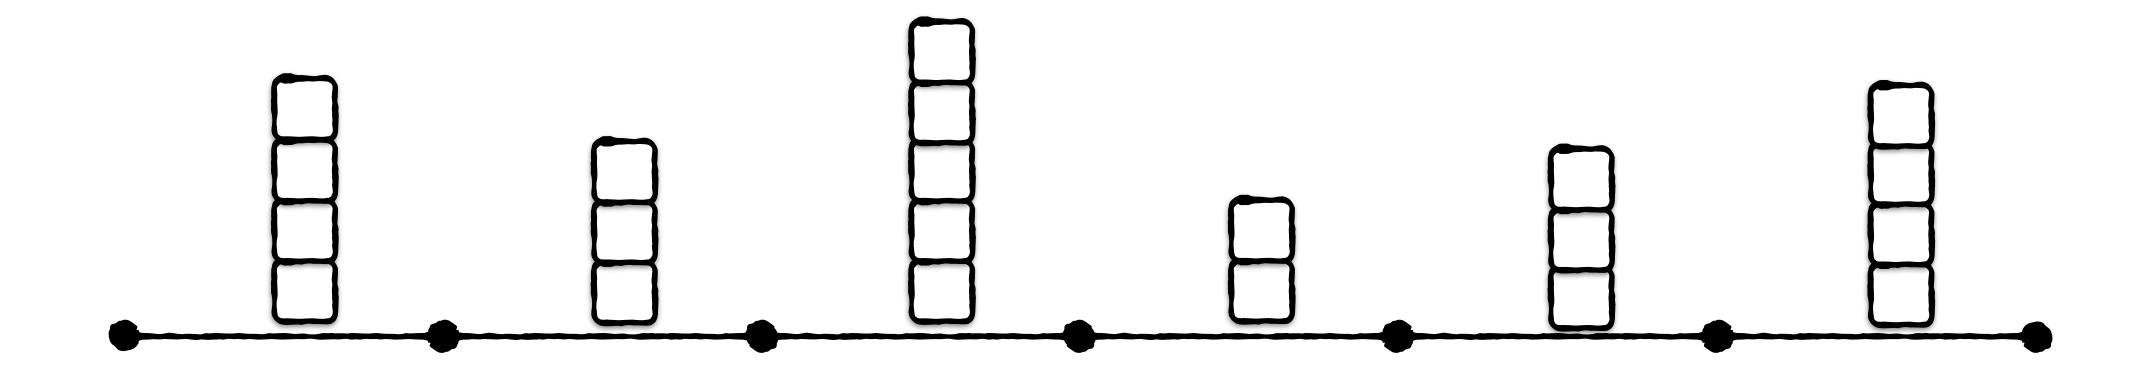
λ = avg. entries per block



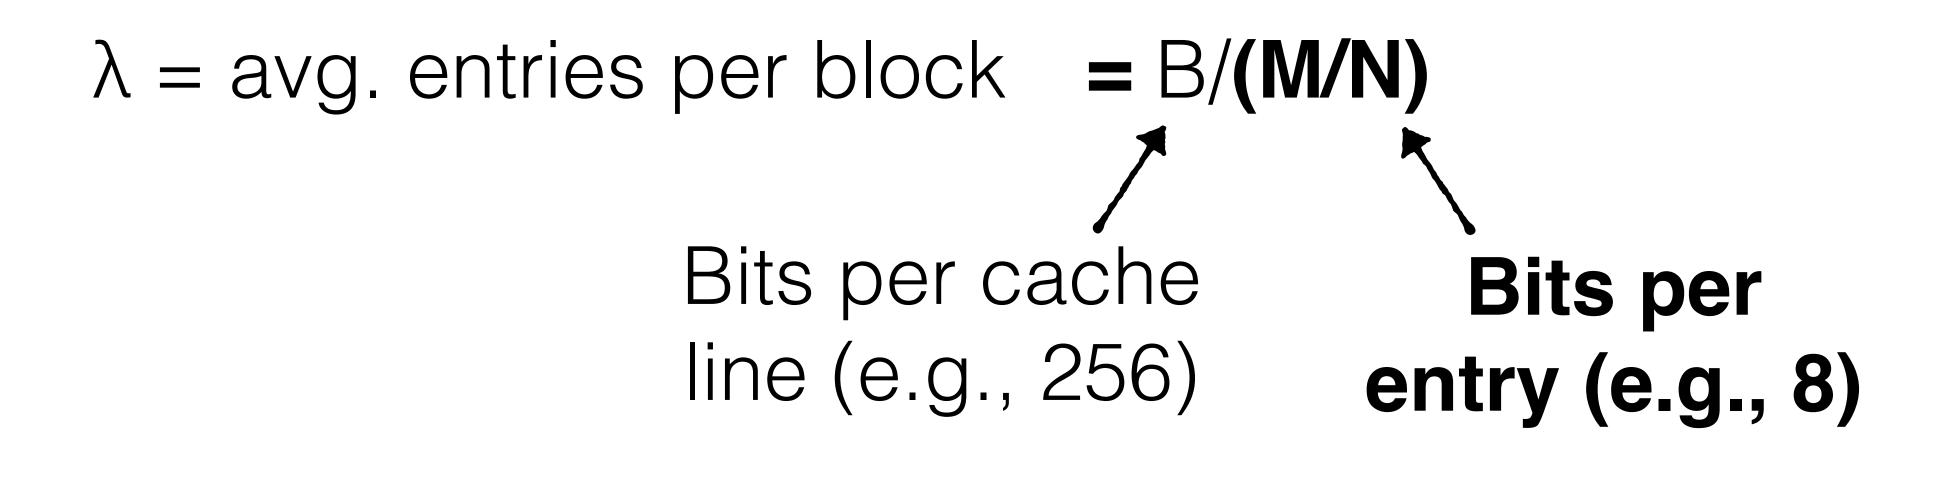
P[i entries fall into given block] ~ Poisson(i,  $\lambda$ ) =  $\frac{\lambda^i \cdot e^{-\lambda}}{i!}$ 

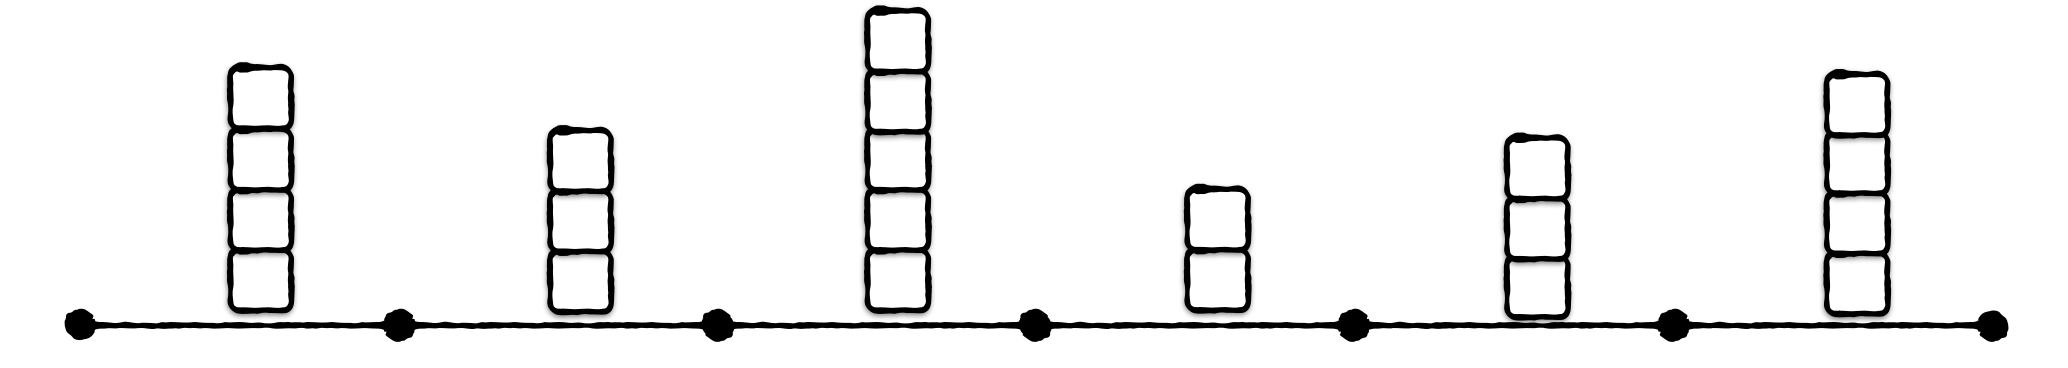
 $\lambda = avg. entries per block = B/(M/N)$ 

Bits per cache line (e.g., 256)



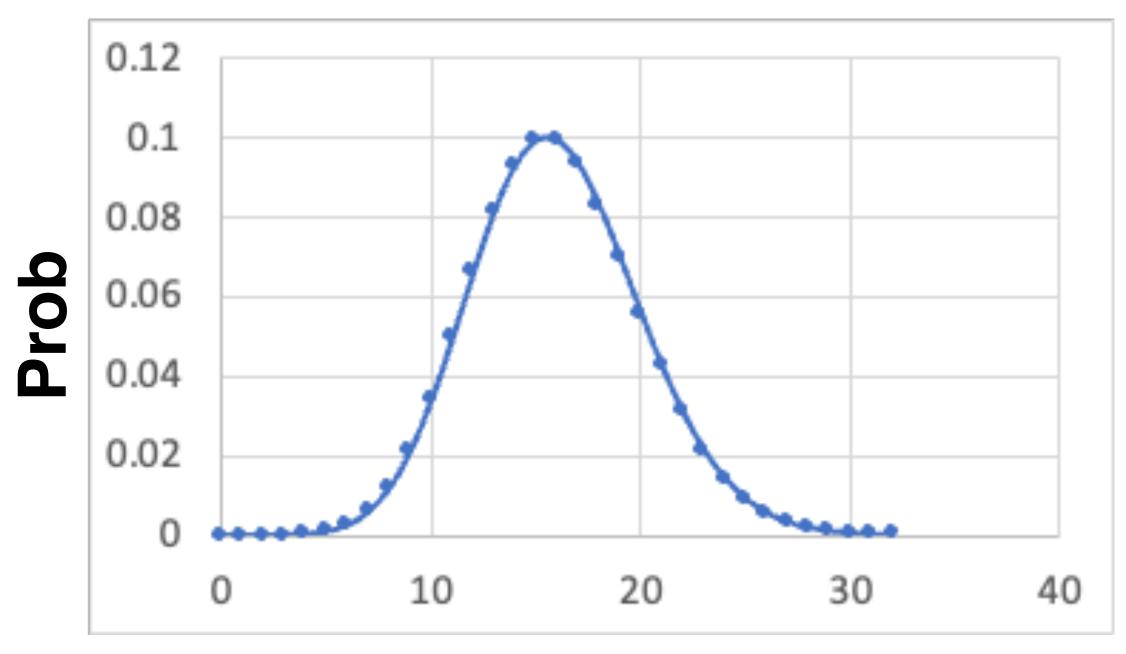
P[i entries fall into given block] ~ Poisson(i,  $\lambda$ ) =  $\frac{\lambda^i \cdot e^{-\lambda}}{i!}$ 





P[i entries fall into given block] ~ Poisson(i,  $\lambda$ ) =  $\frac{\lambda^{i} \cdot e^{-\lambda}}{i!}$ 

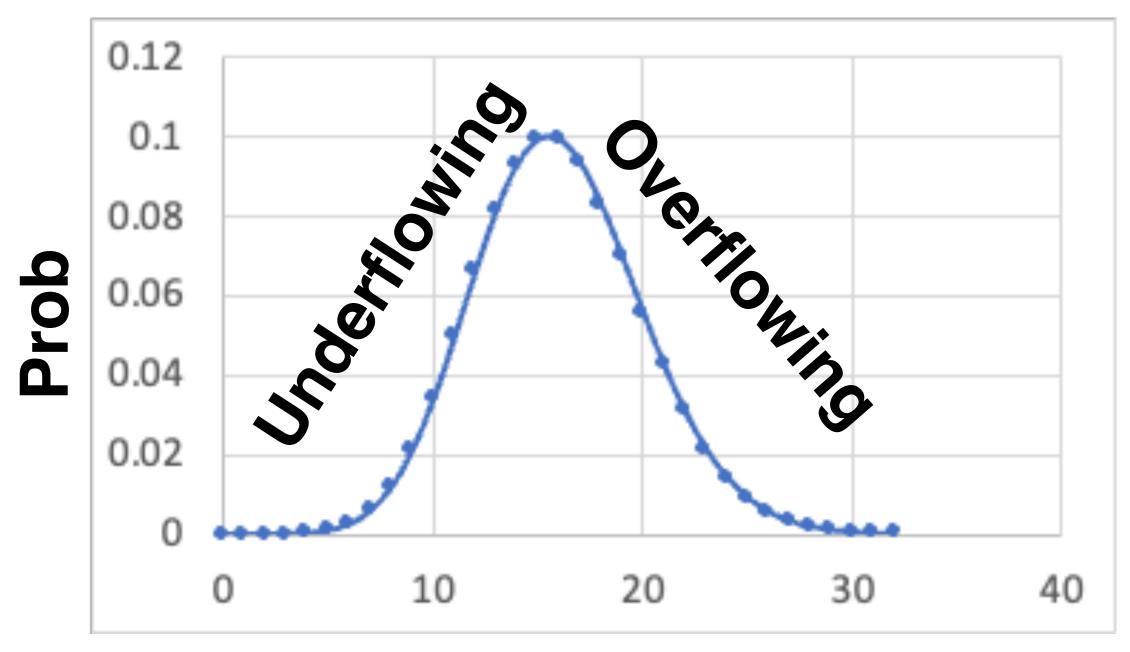
 $\lambda$  = avg. entries per block



Entries i per block with  $\lambda = 16$ 

P[i entries fall into given block] ~ Poisson(i,  $\lambda$ ) =  $\frac{\lambda^{i} \cdot e^{-\lambda}}{i!}$ 

 $\lambda$  = avg. entries per block



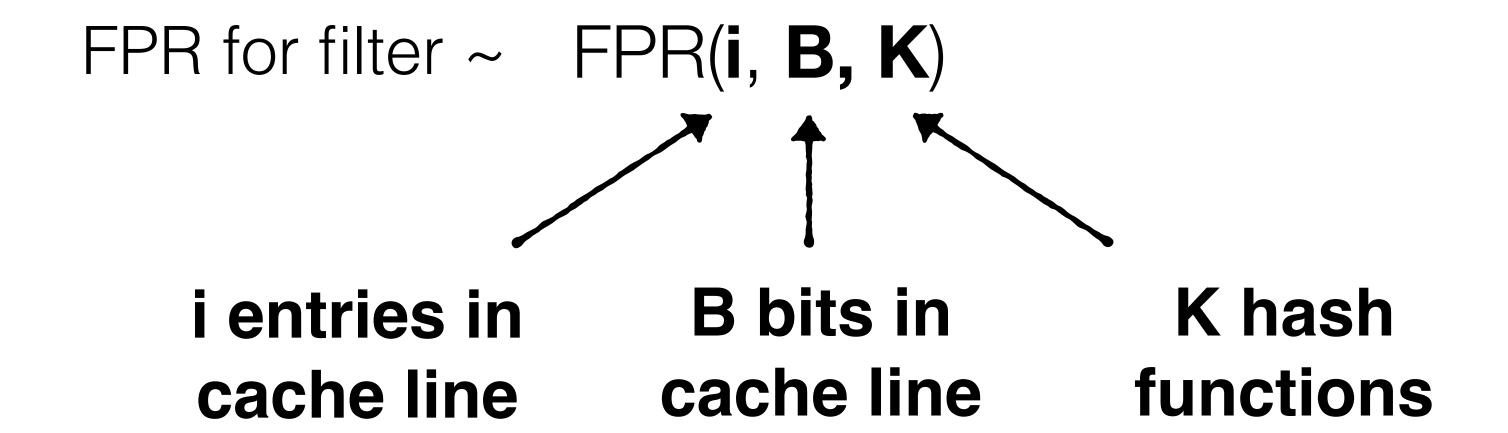
Entries i per block with  $\lambda = 16$ 

# entries in a block ~ Poisson(i, B/(M/N))

# entries in a block ~ Poisson(i, B/(M/N))

FPR for filter ~ FPR(N, M, K) =  $(1-e^{-KN/M})^{K}$ 

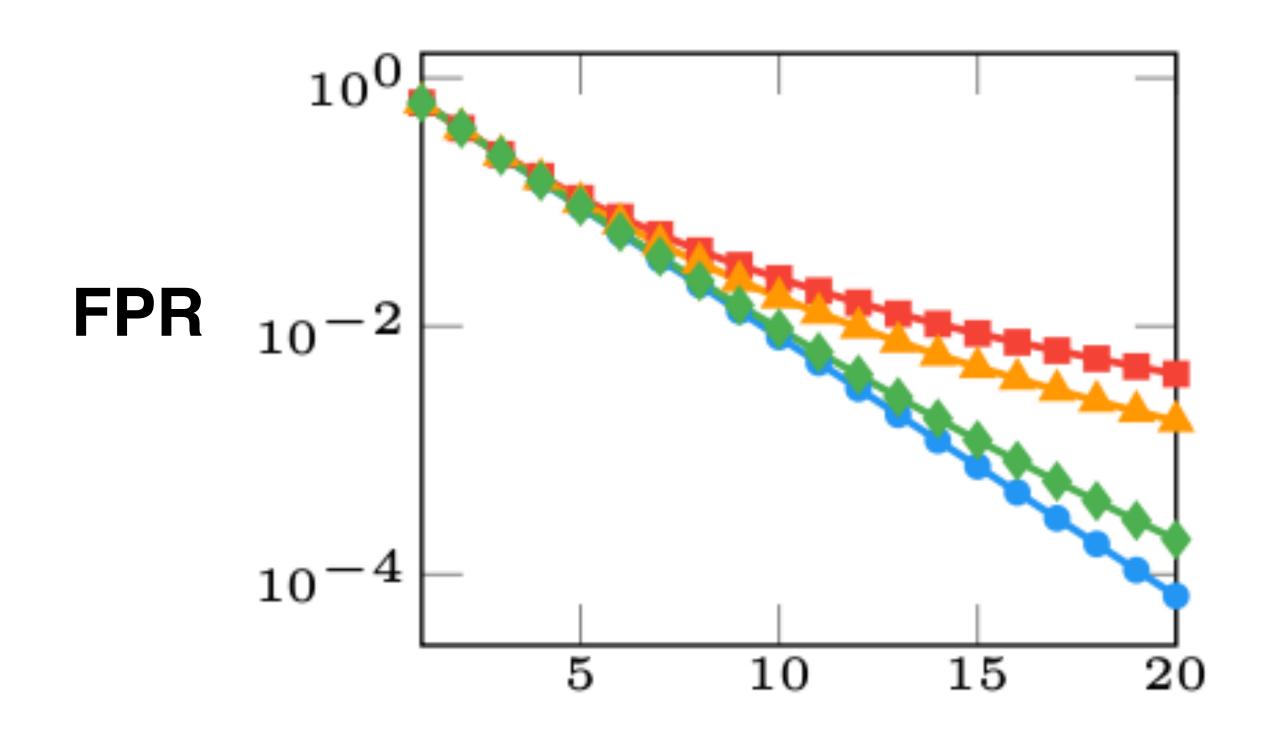
# entries in a block ~ Poisson(i, B/(M/N))



Avg. FPR across all blocks

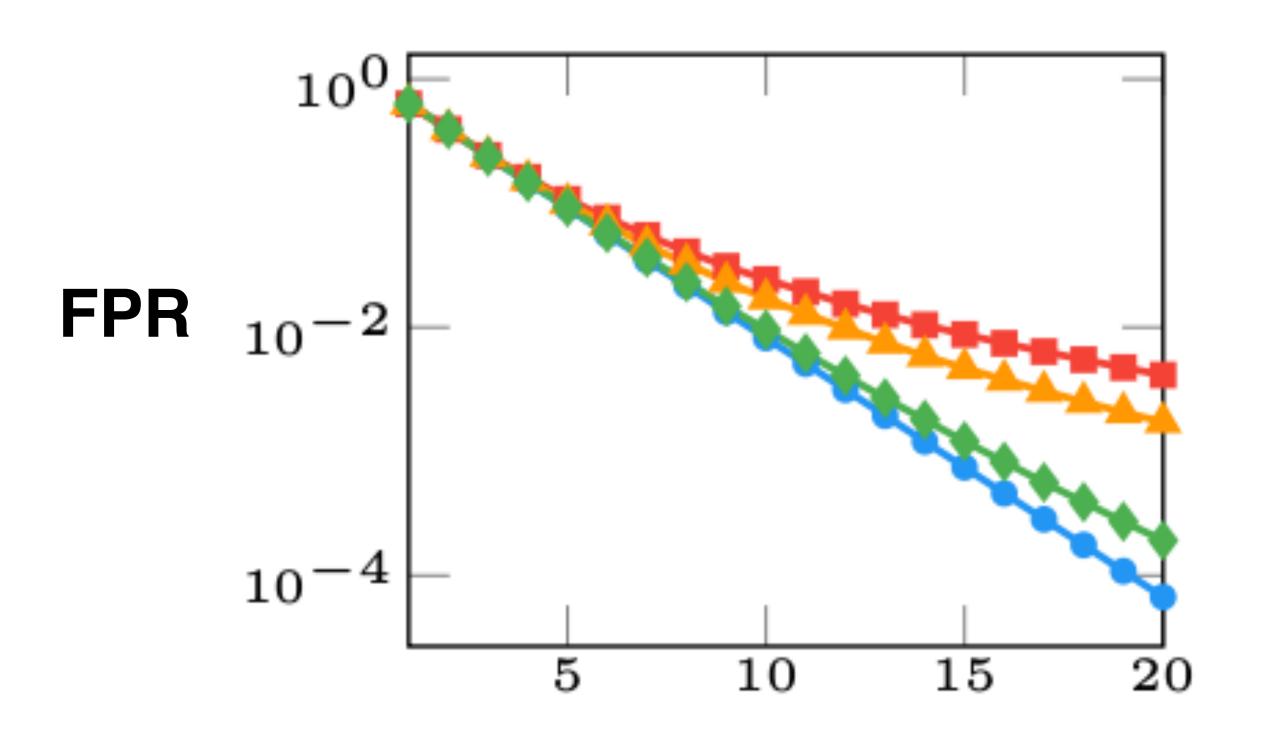
 $\infty$ 

Classic Bloom
32-bit blocked
64-bit blocked
512-bit blocked



Bits / entry (M/N)

Classic Bloom
32-bit blocked
64-bit blocked
512-bit blocked

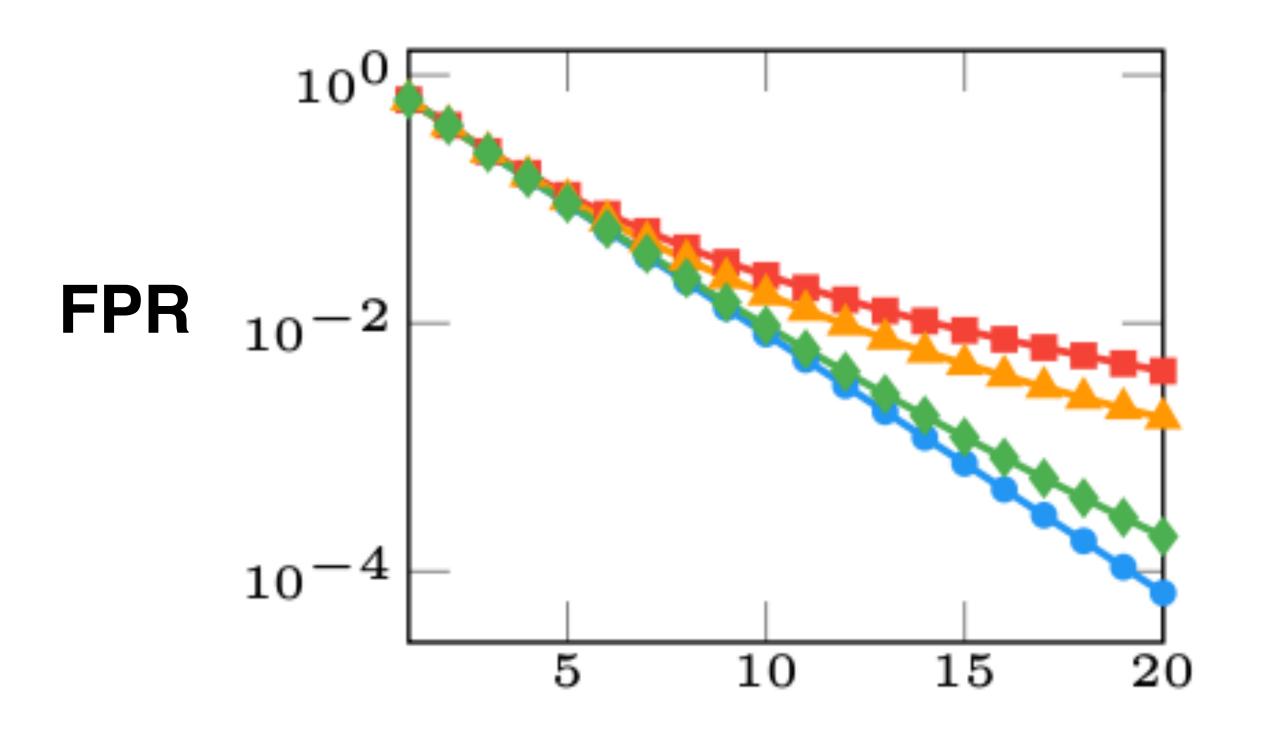


With smaller blocks, there is more variation in entries across blocks.

Overflowing blocks blow up the FPR.

Bits / entry (M/N)

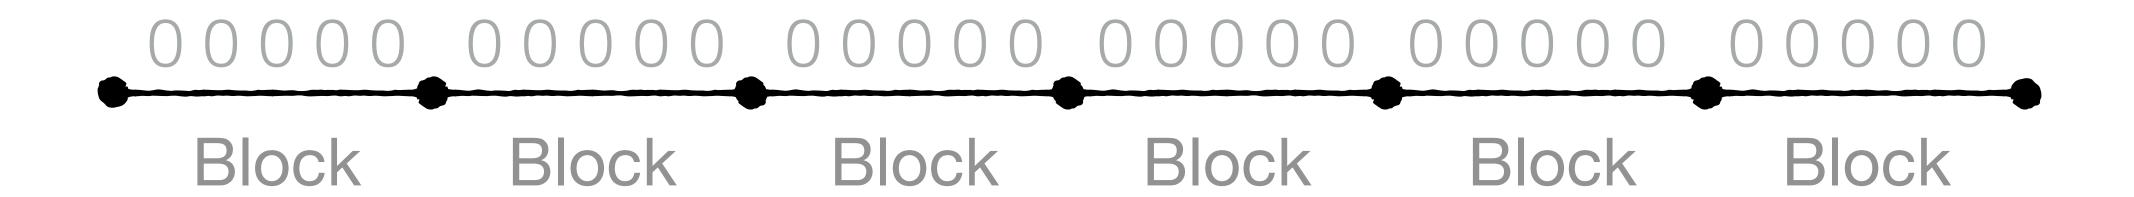
Classic Bloom
32-bit blocked
64-bit blocked
512-bit blocked



With a block the size of a cache line, we don't lose much

Bits / entry (M/N)

#### size blocks as cache lines



#### size blocks as cache lines

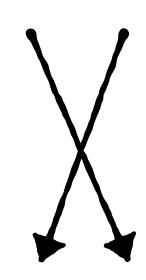
	00000				00000
Block	Block	Block	Block	Block	Block
512	512	512	512	512	512
bits	bits	bits	bits	bits	bits

### Any remaining issue?

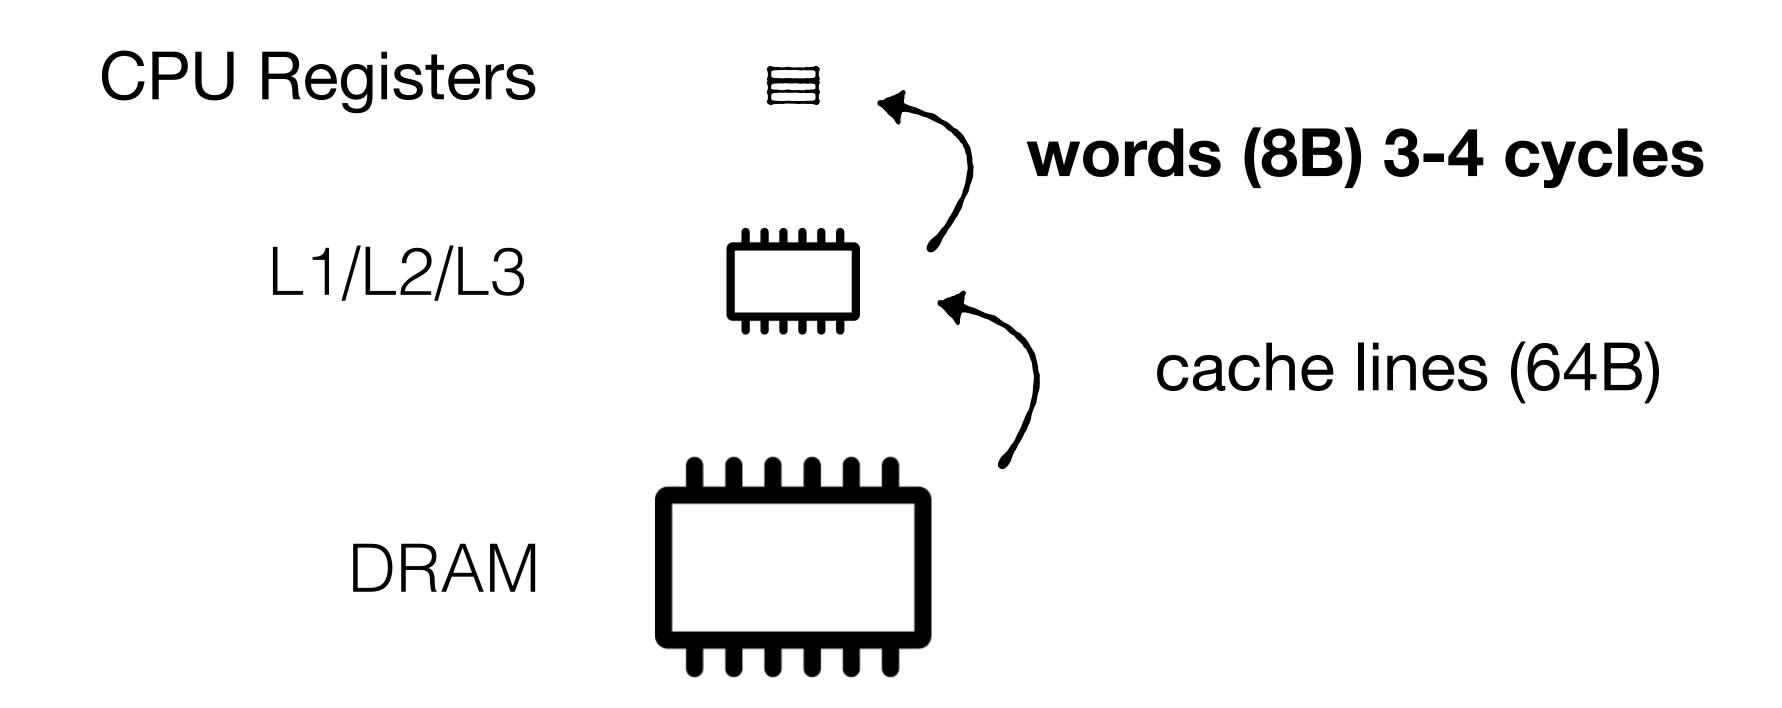
00000	00000					
Block	Block	Block	Block	Block	Block	
512	512	512	512	512	512	
bits	bits	bits	bits	bits	bits	

### Any remaining issue?

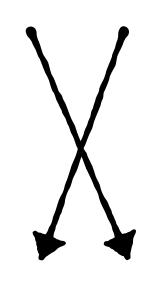
#### Random access within a cache line

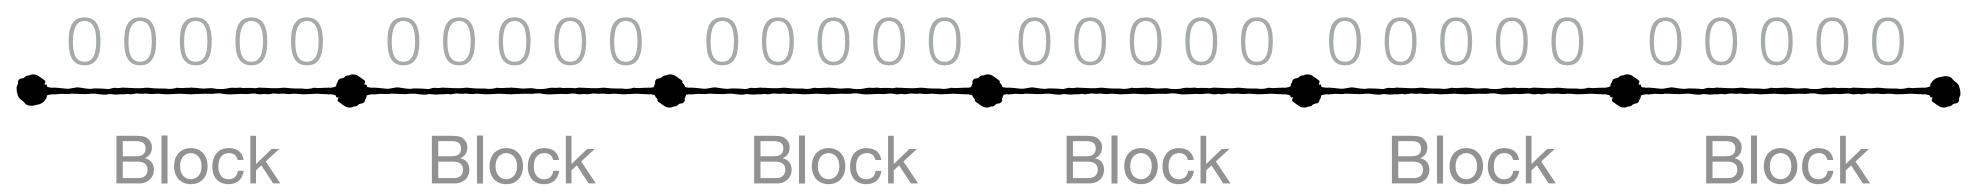


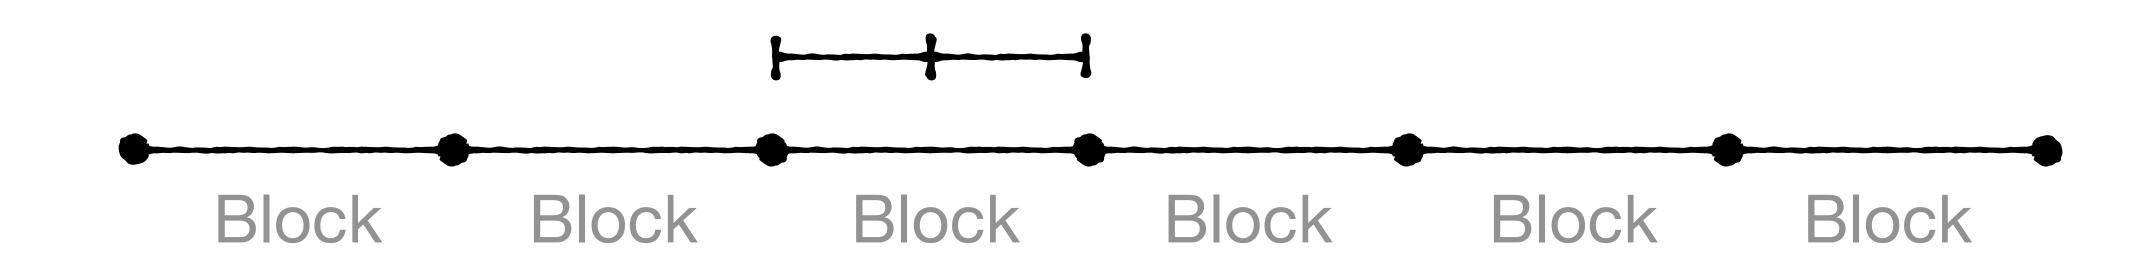
	00000					
Block	Block	Block	Block	Block	Block	
512	512	512	512	512	512	
bits	bits	bits	bits	bits	bits	

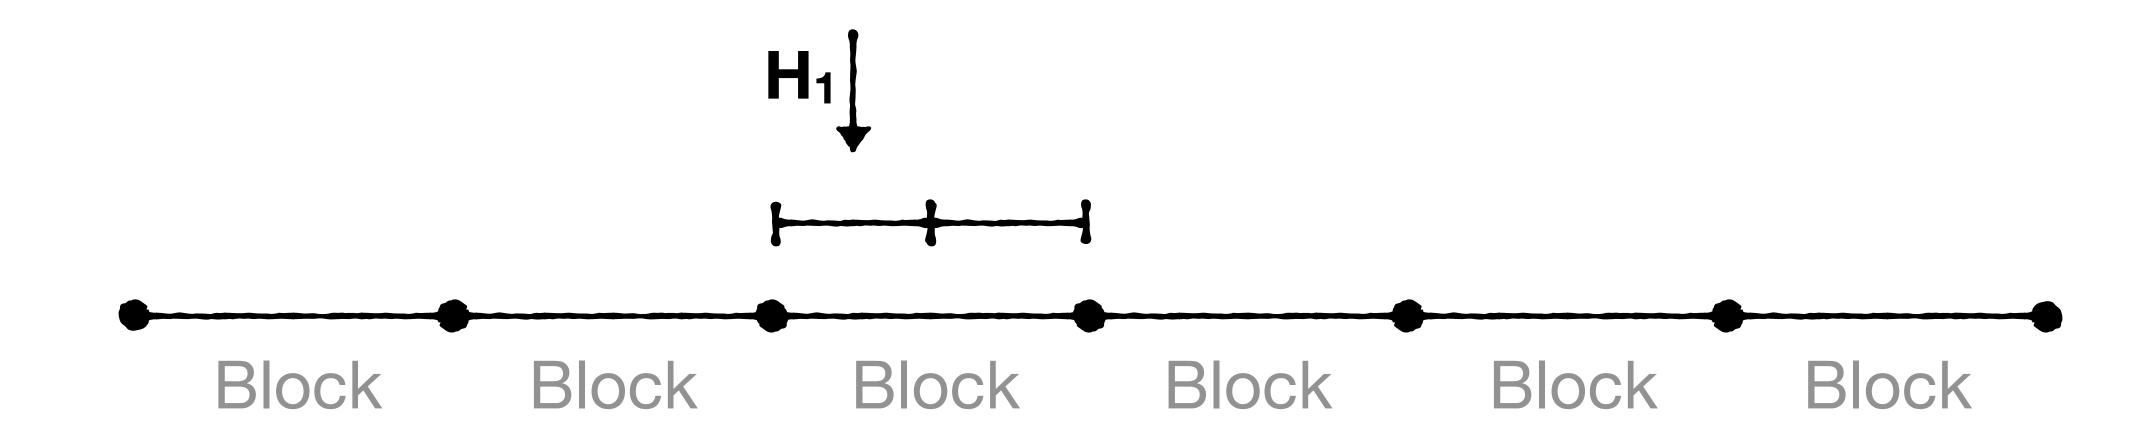


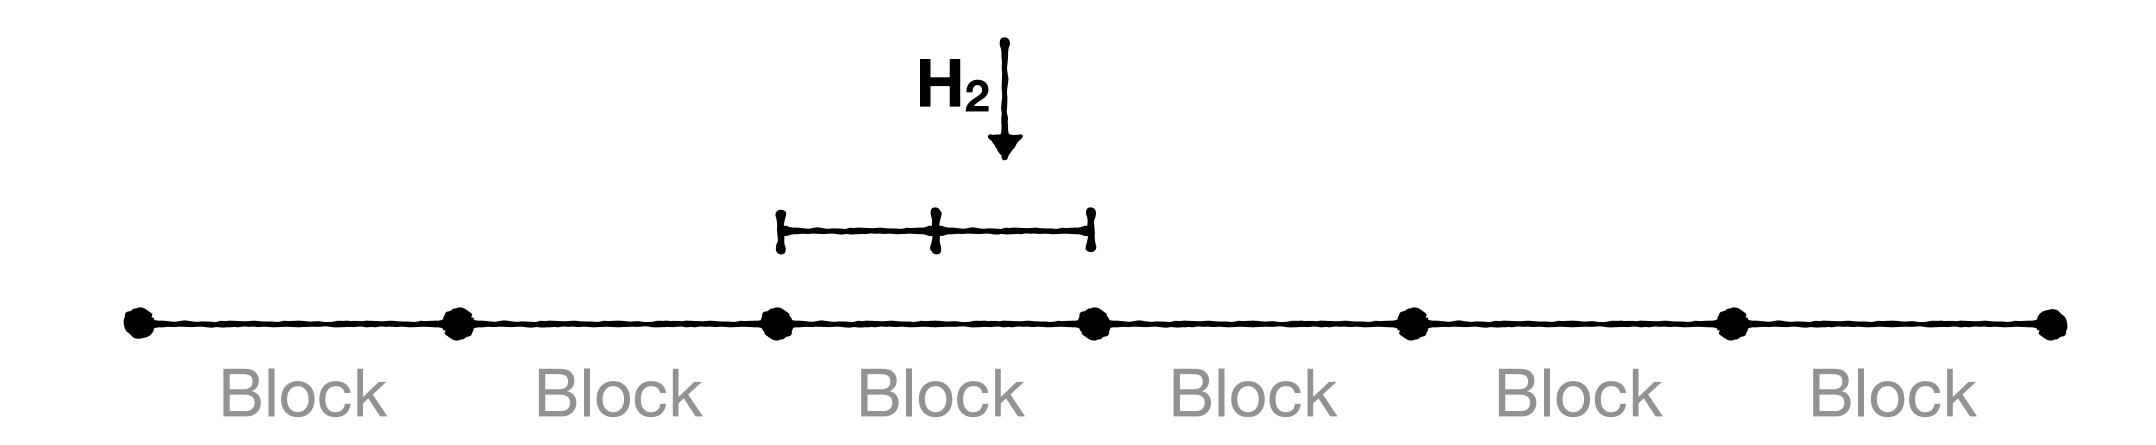
# Each random access moves different word from L1 cache to register

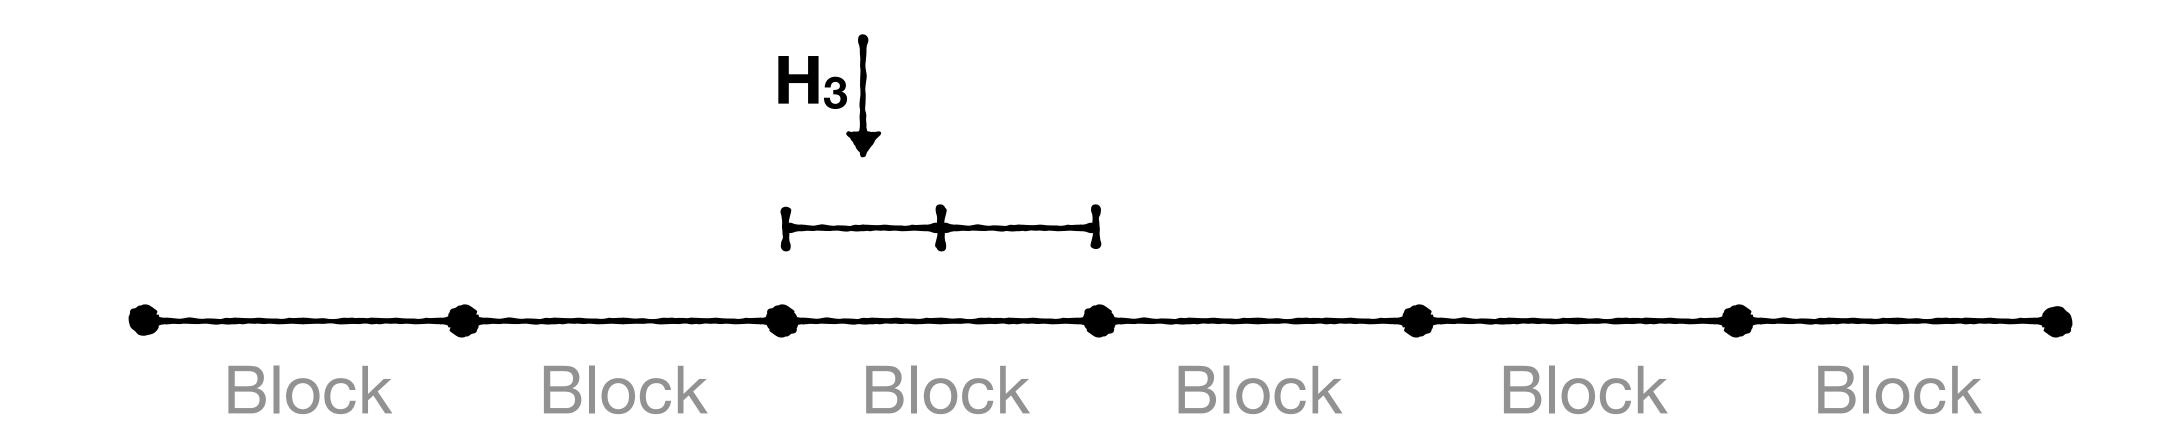




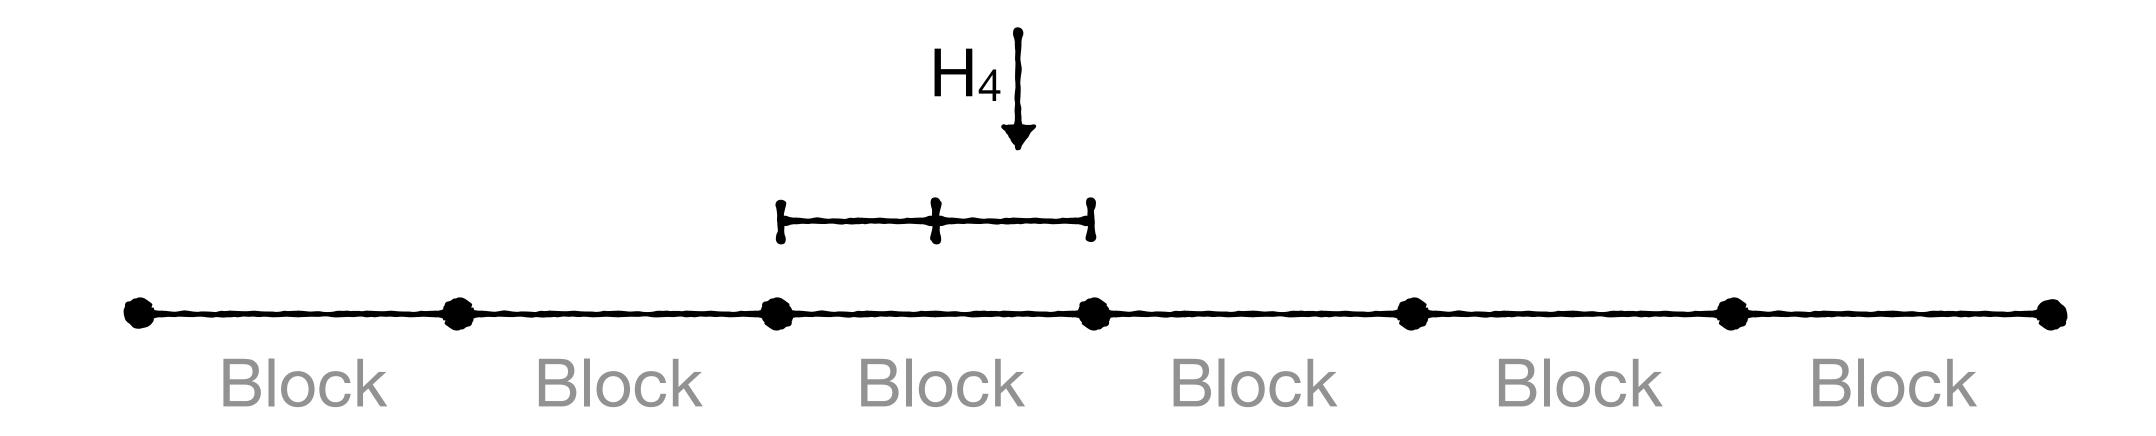








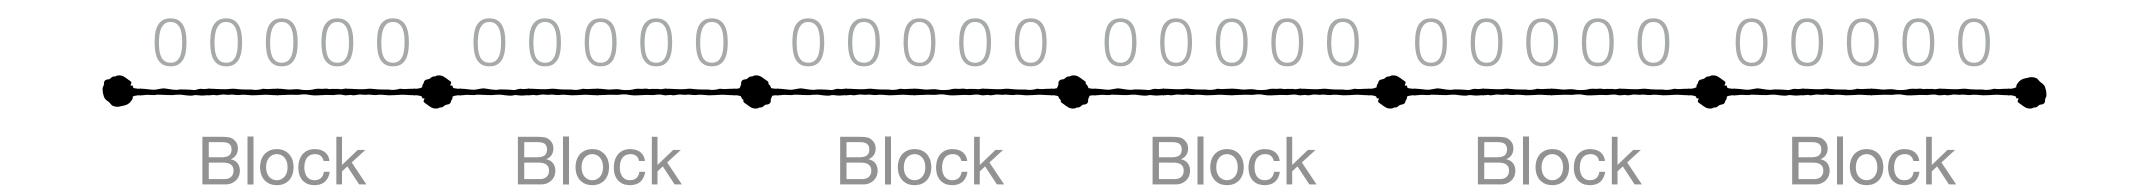
#### Solutions?



#### Split block Bloom filters. Jim Apple. Arxiv 2023.

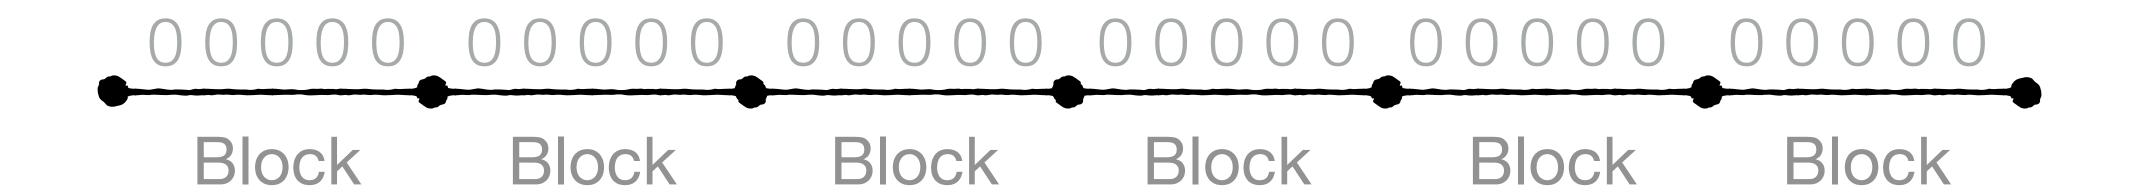
Used in the Impala system as of 2016



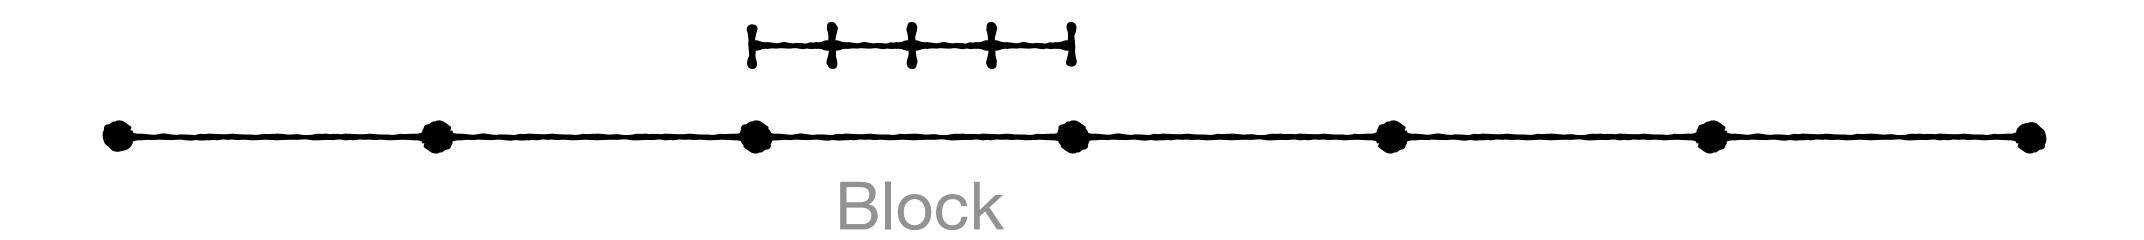


Split block Bloom filters. Jim Apple. Arxiv 2023.

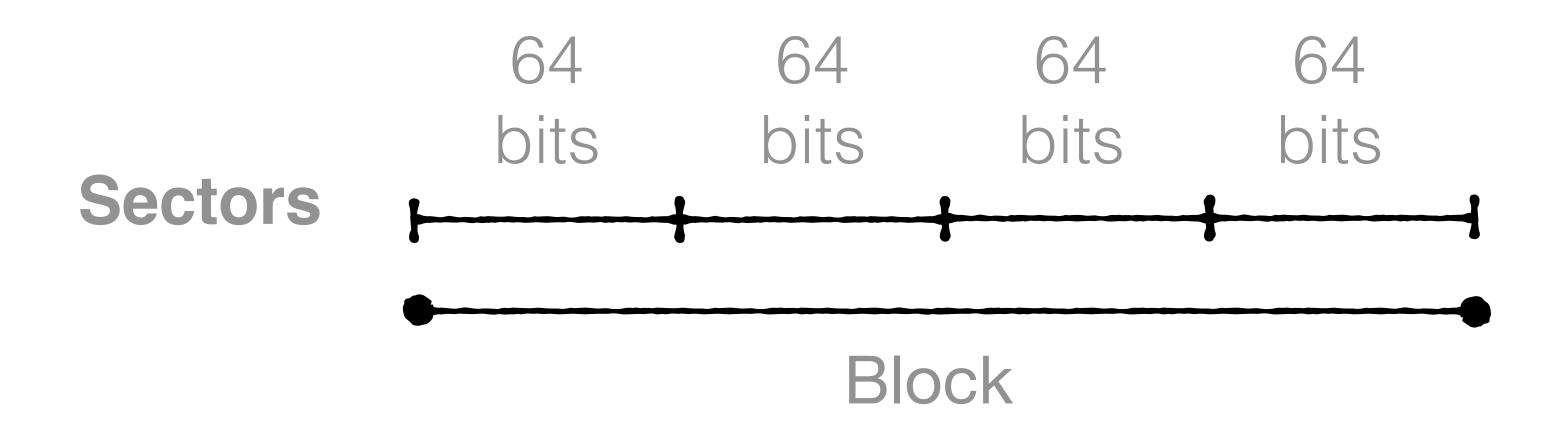
Performance-Optimal Filtering: Bloom Overtakes Cuckoo at High Throughput. Harald Lang, Thomas Neumann, Alfons Kemper, Peter Boncz. VLDB 2019.

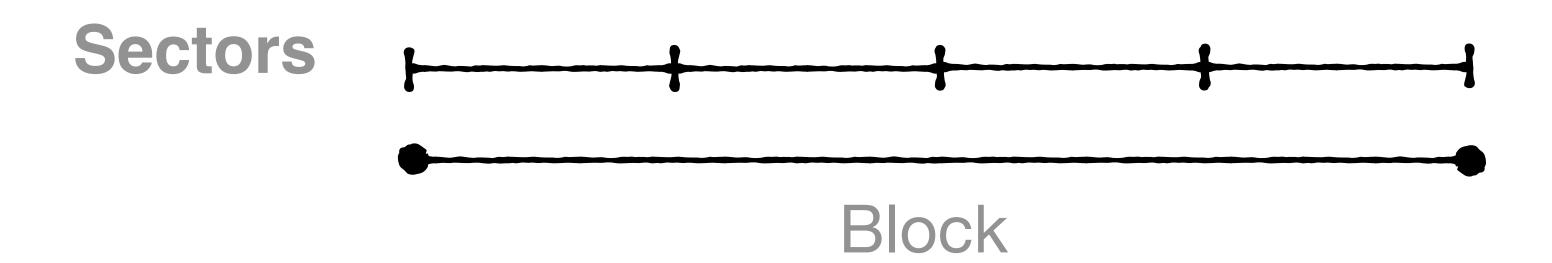


# Partition into s sectors, each the size of a word (e.g., 64 bits)

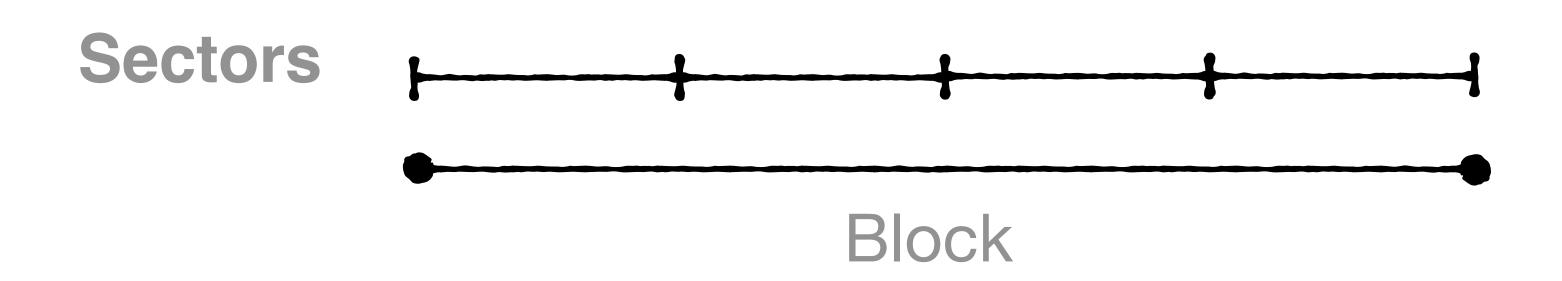


## Example: s = 4 sectors

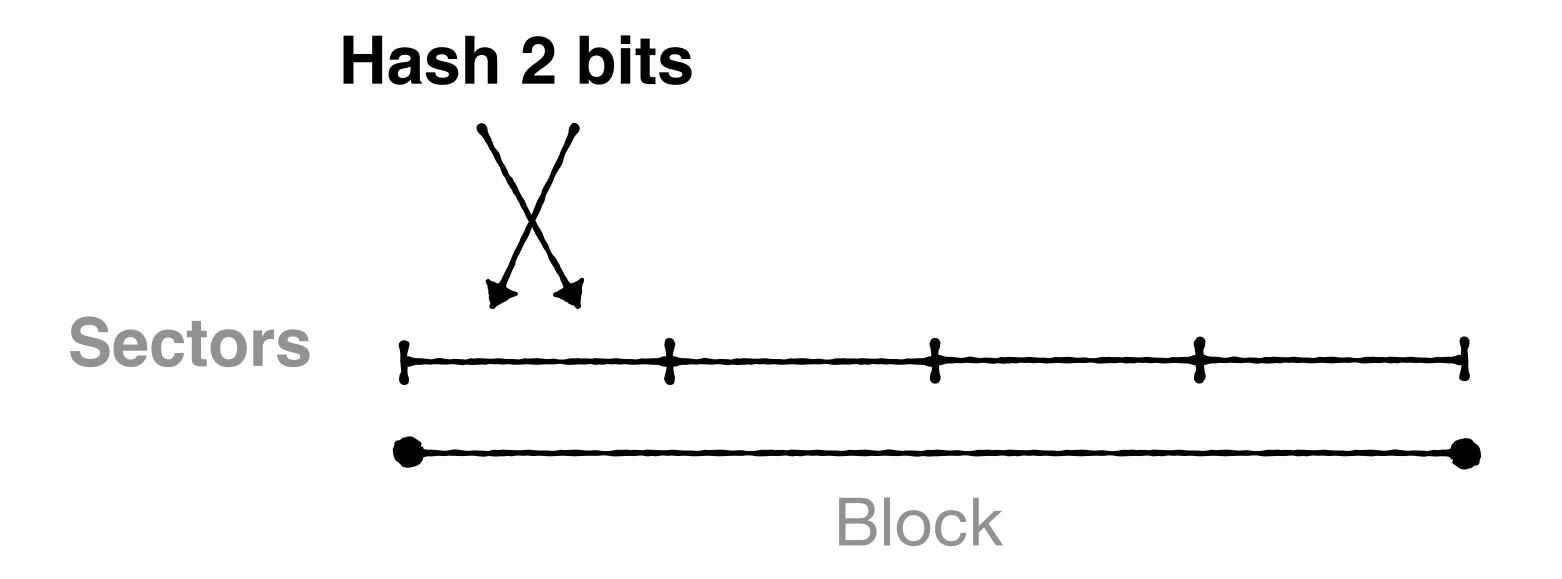




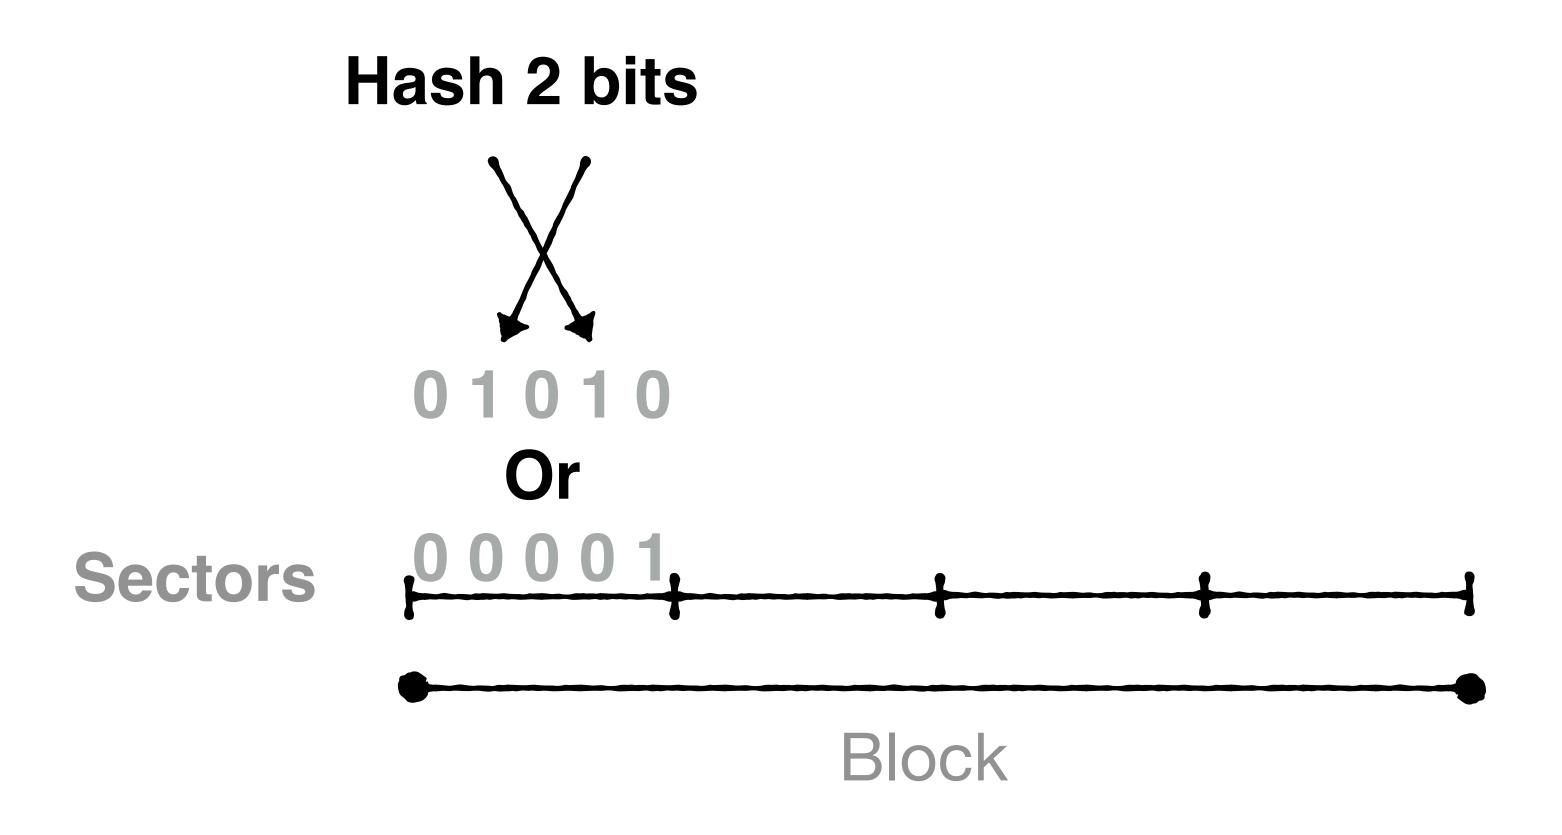
e.g., s=4 sectors and k=8 hashes



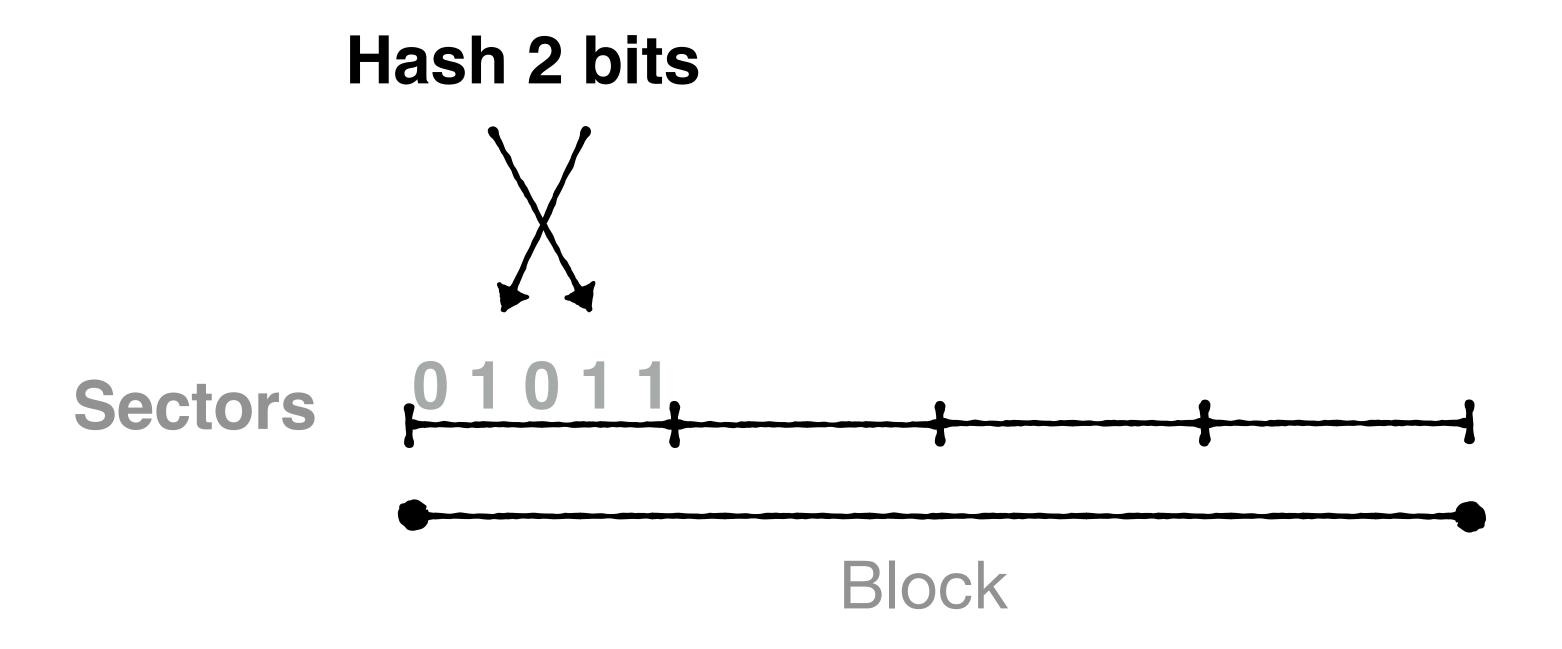
e.g., s=4 sectors and k=8 hashes



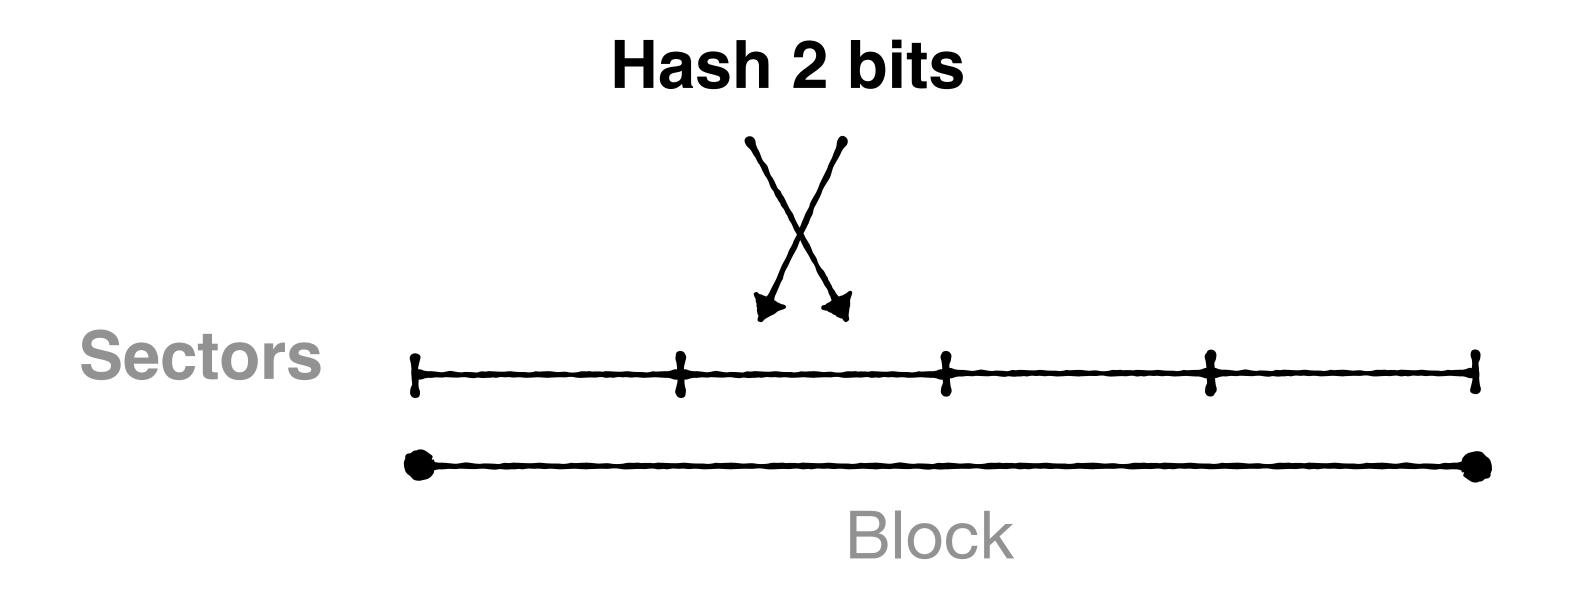
e.g., s=4 sectors and k=8 hashes



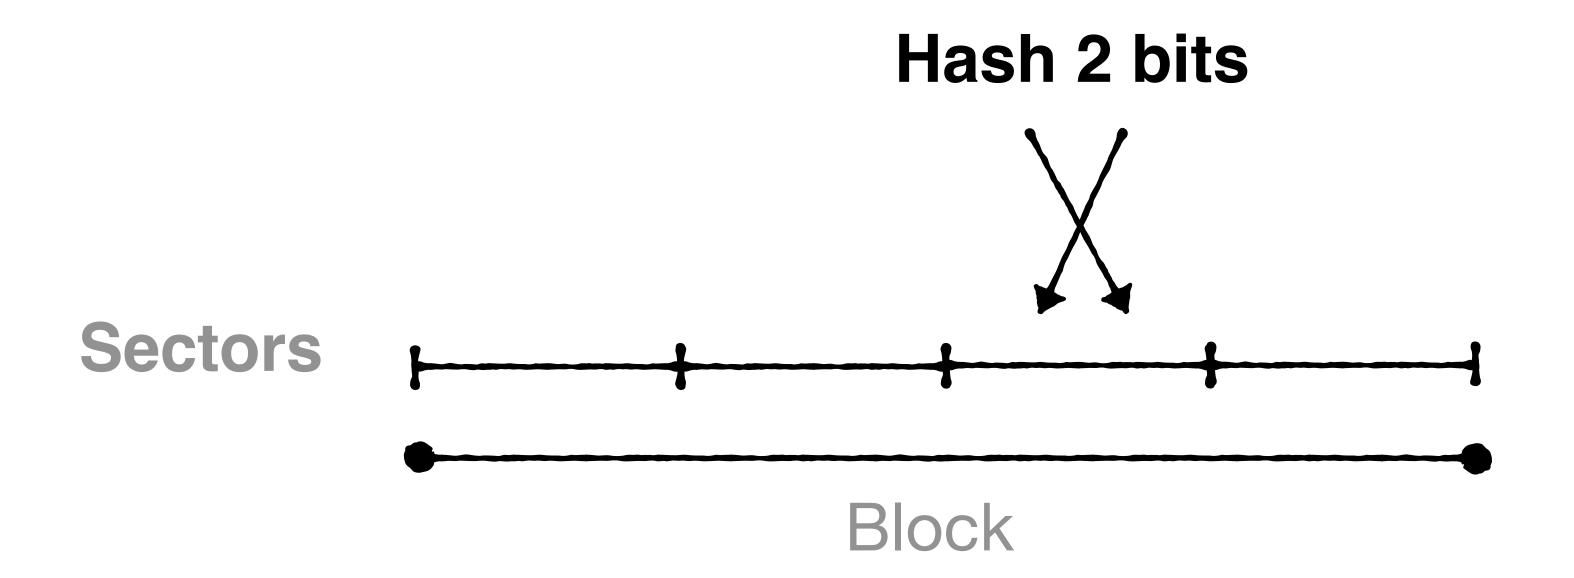
e.g., s=4 sectors and k=8 hashes



e.g., s=4 sectors and k=8 hashes

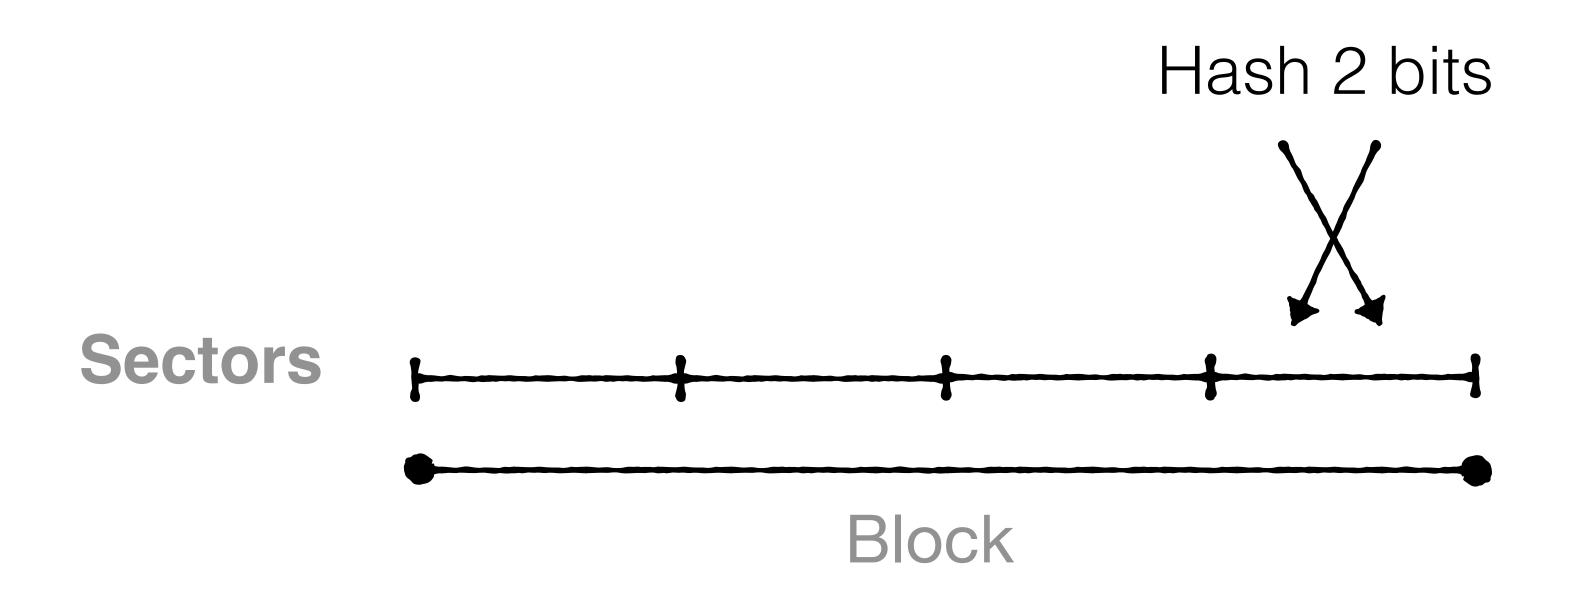


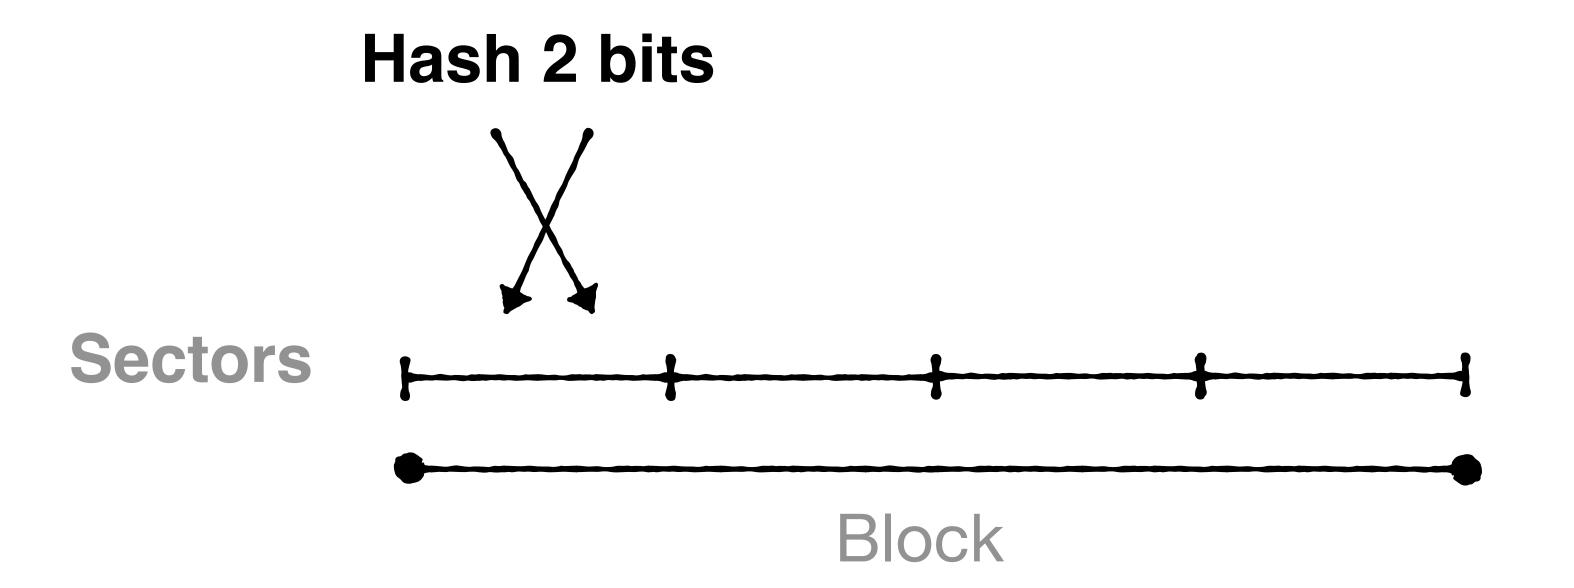
e.g., s=4 sectors and k=8 hashes

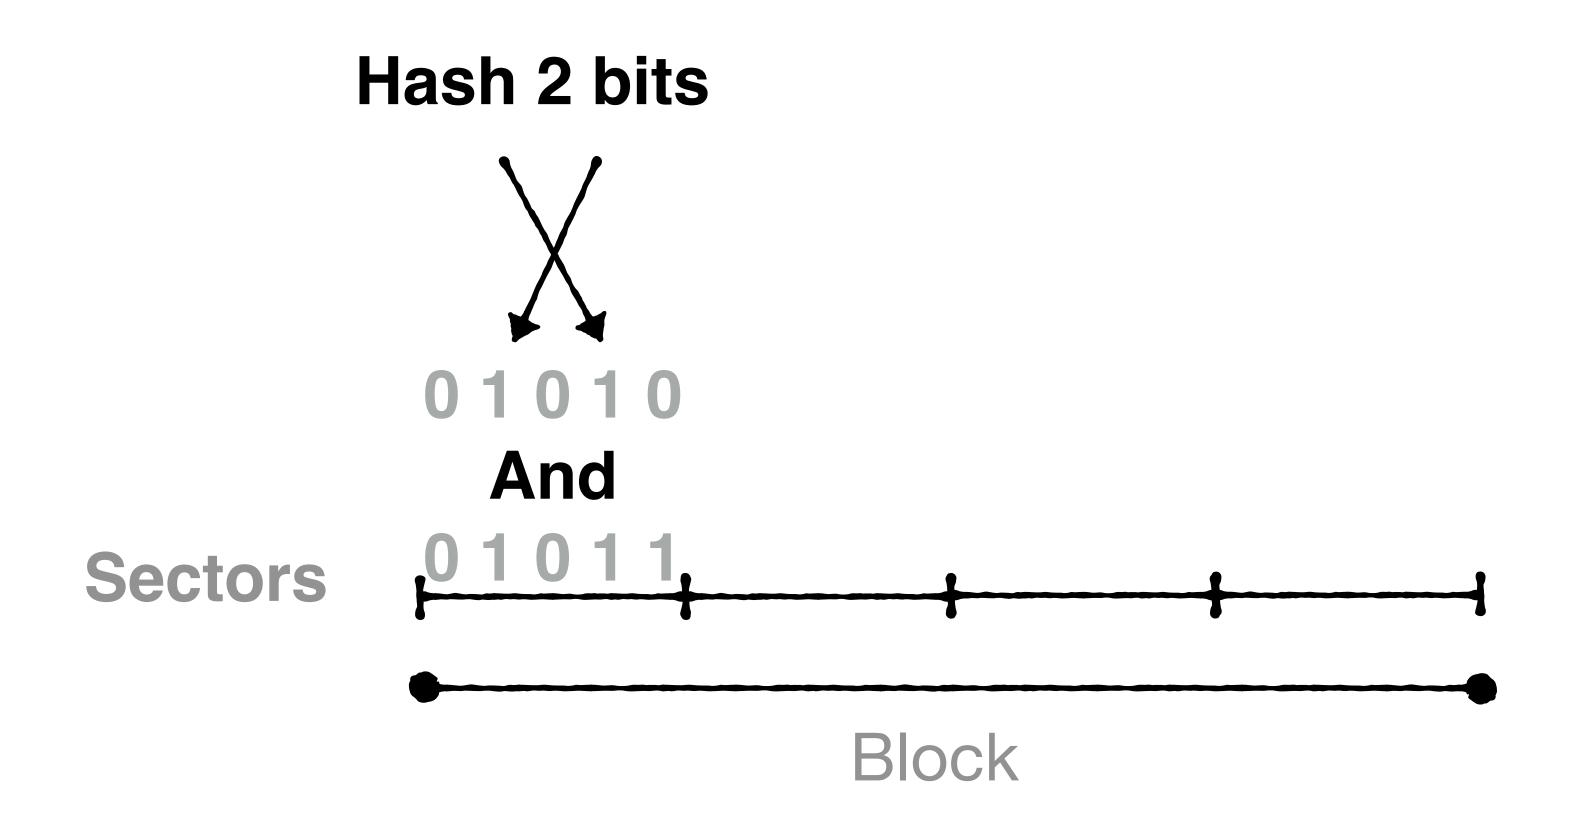


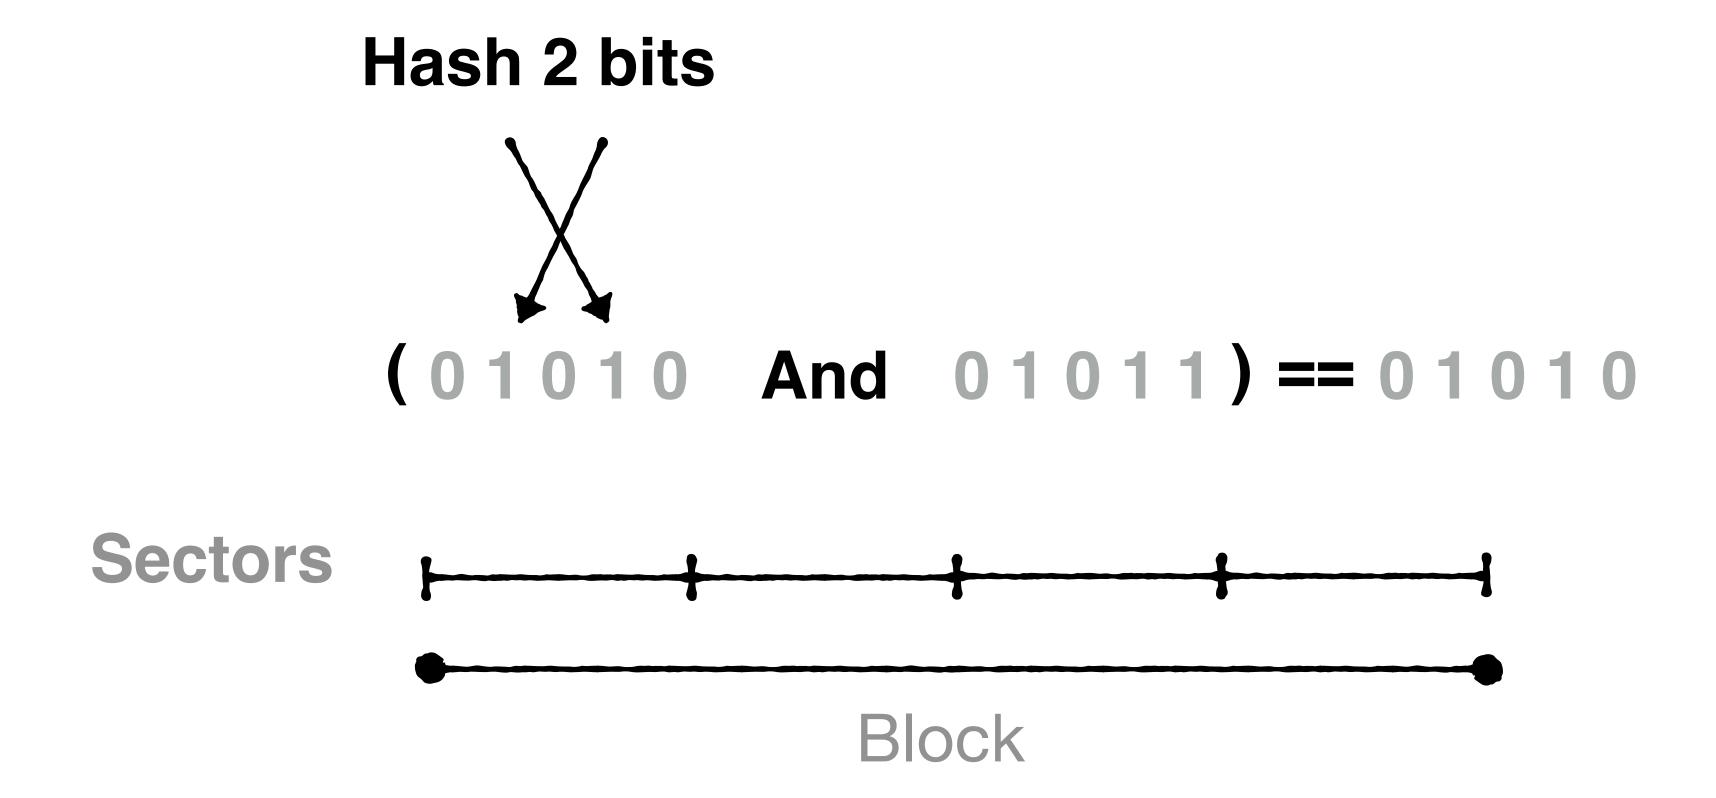
Insertion: hashes to k / s bits per sector sequentially

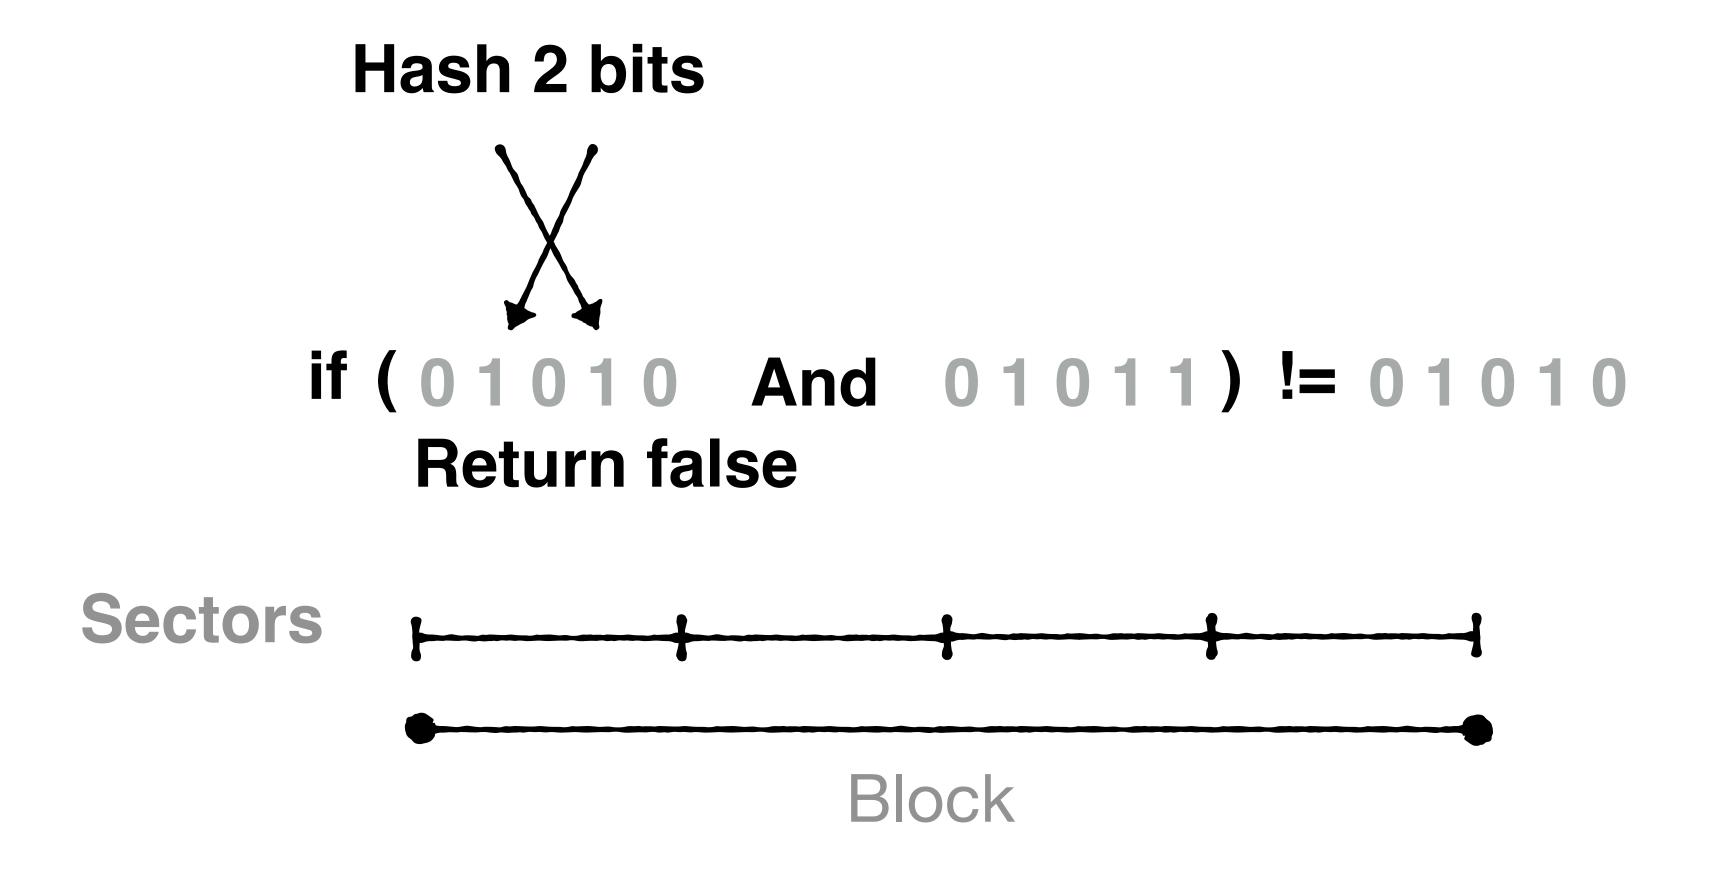
### Read/written 4 rather than 8 registers

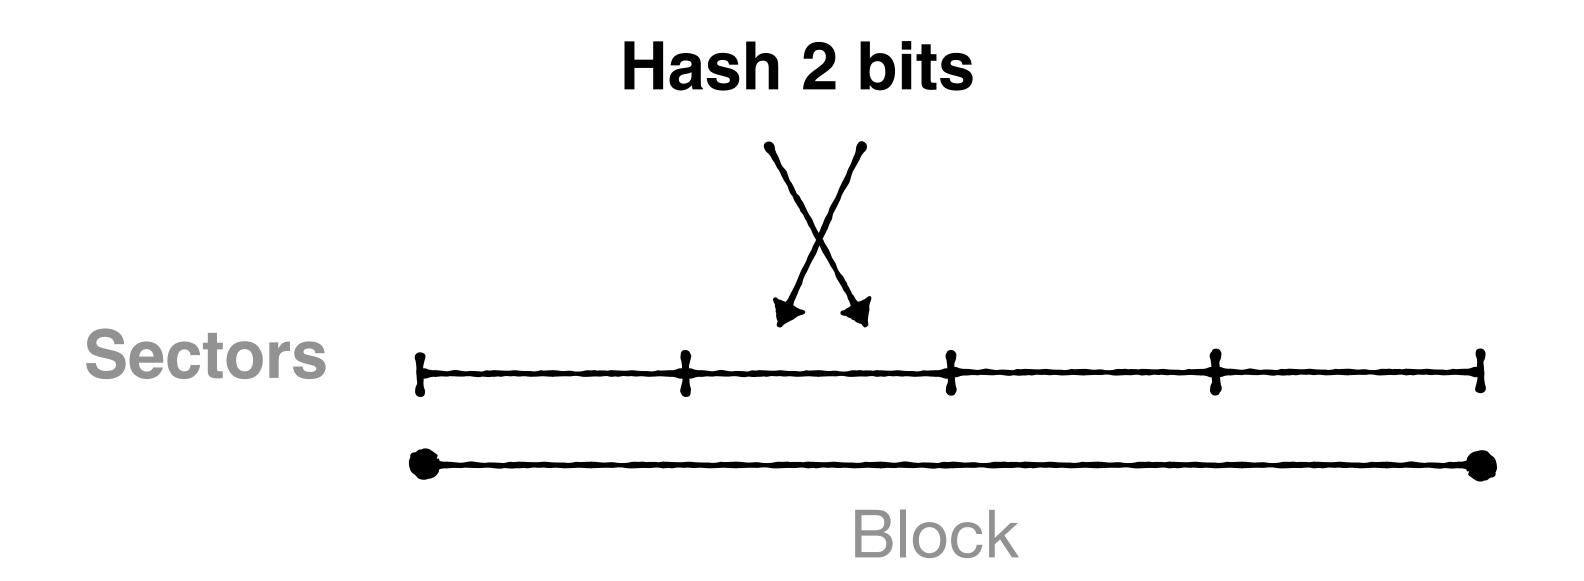


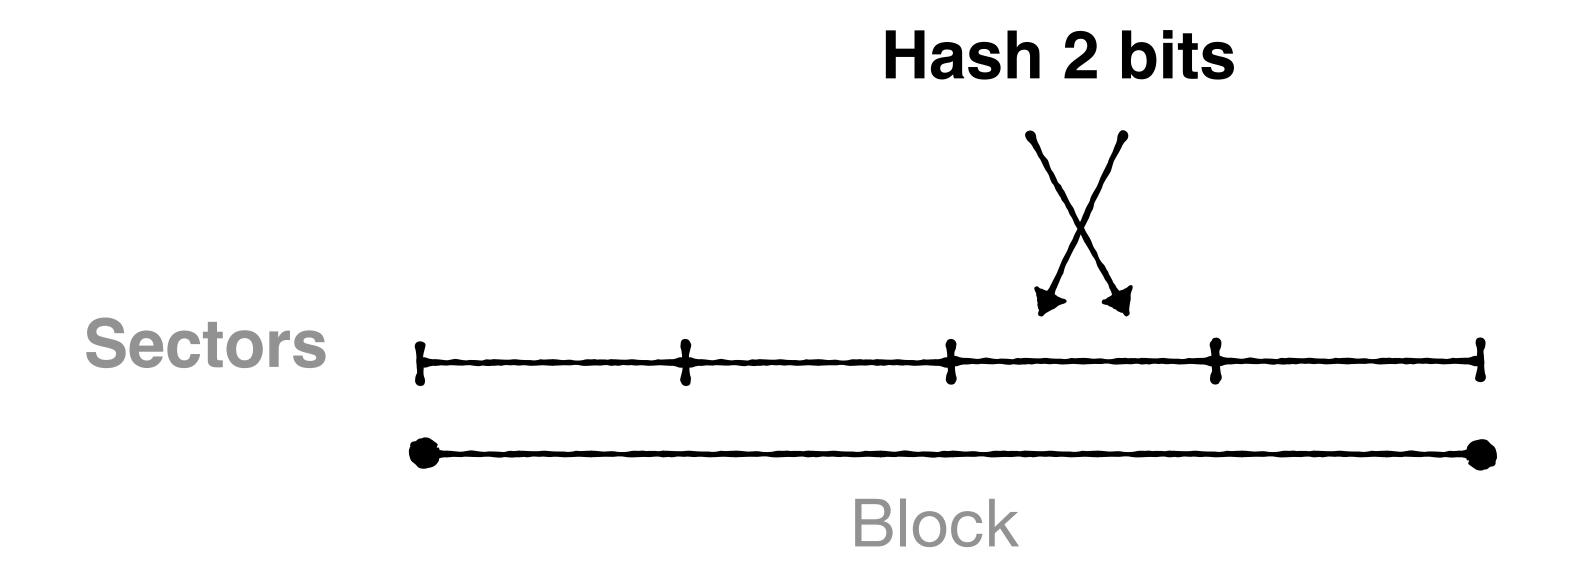


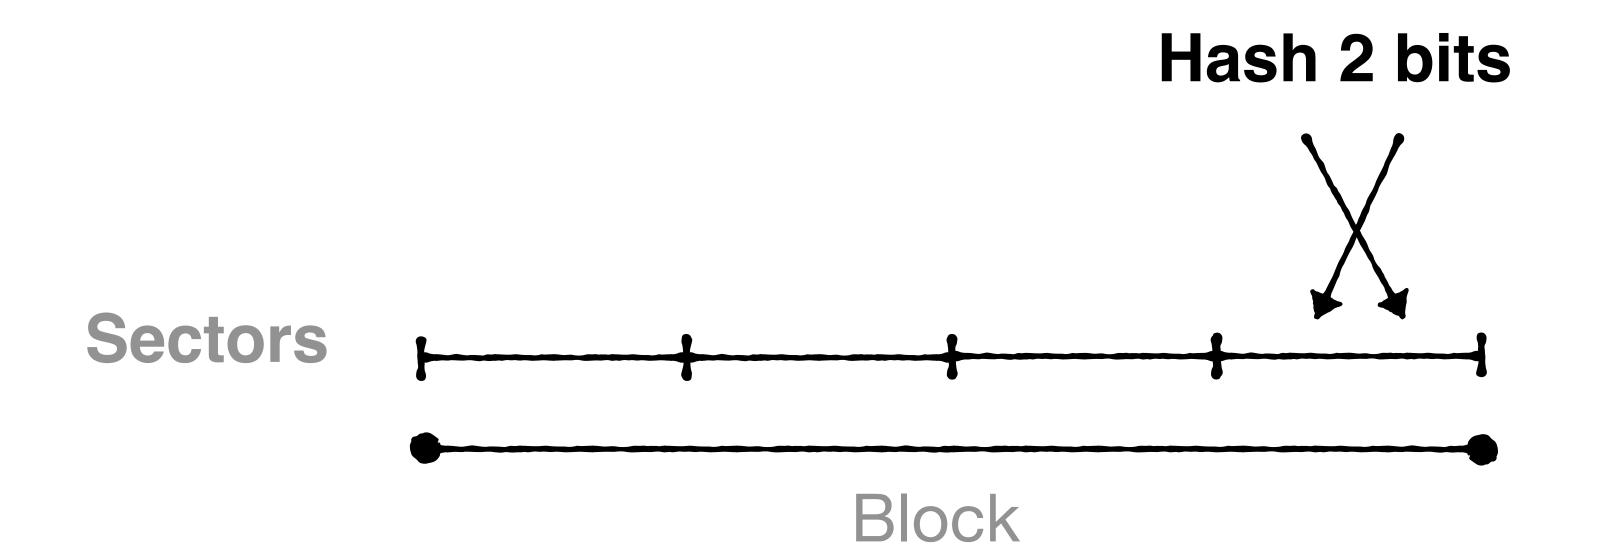


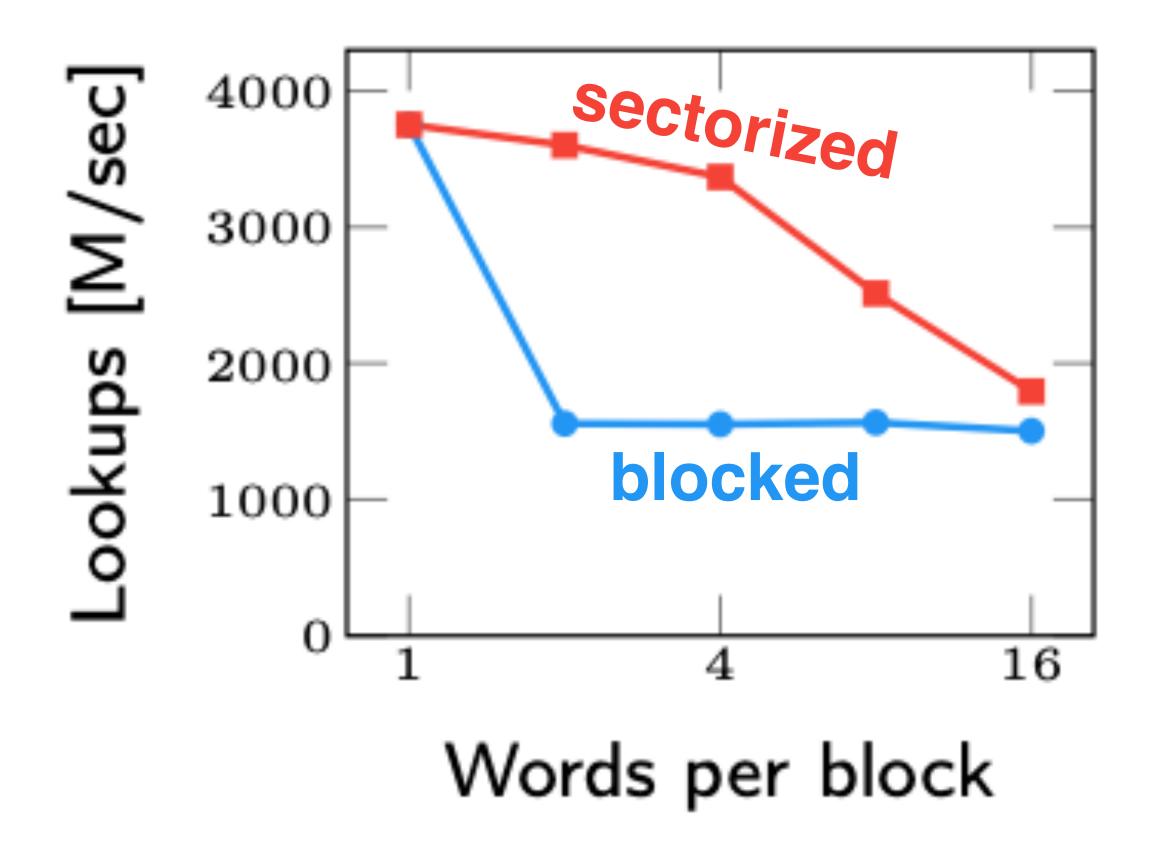




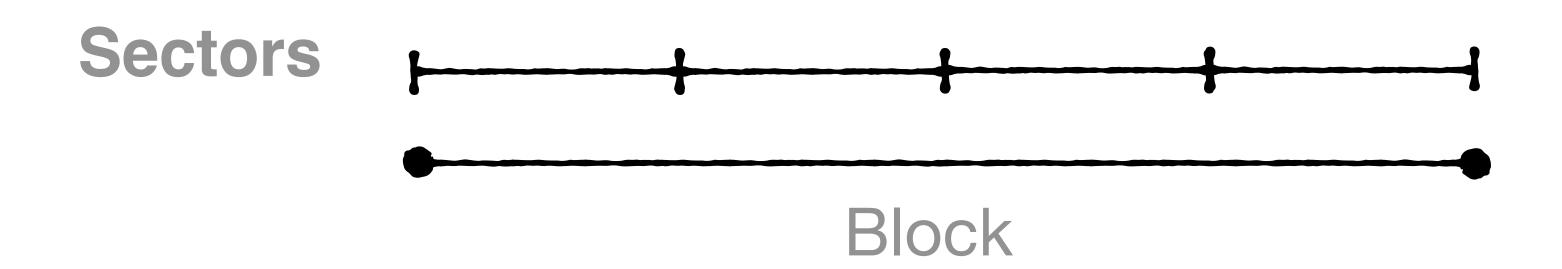




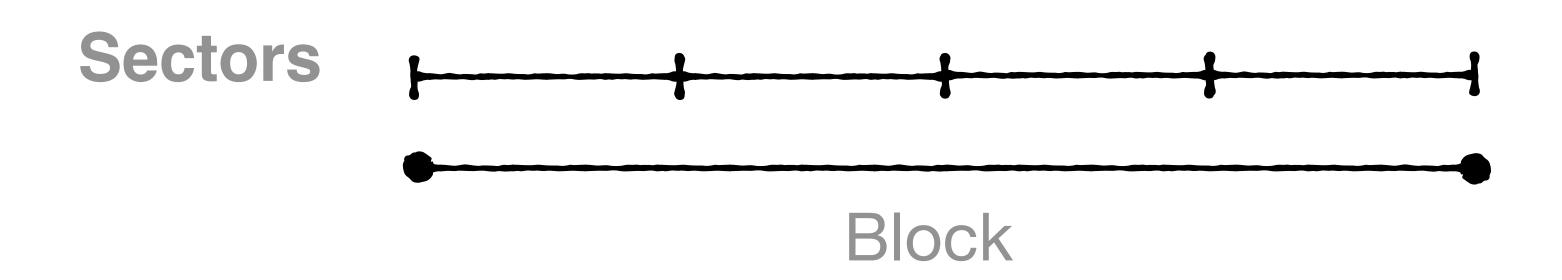




## Further ways to optimize?



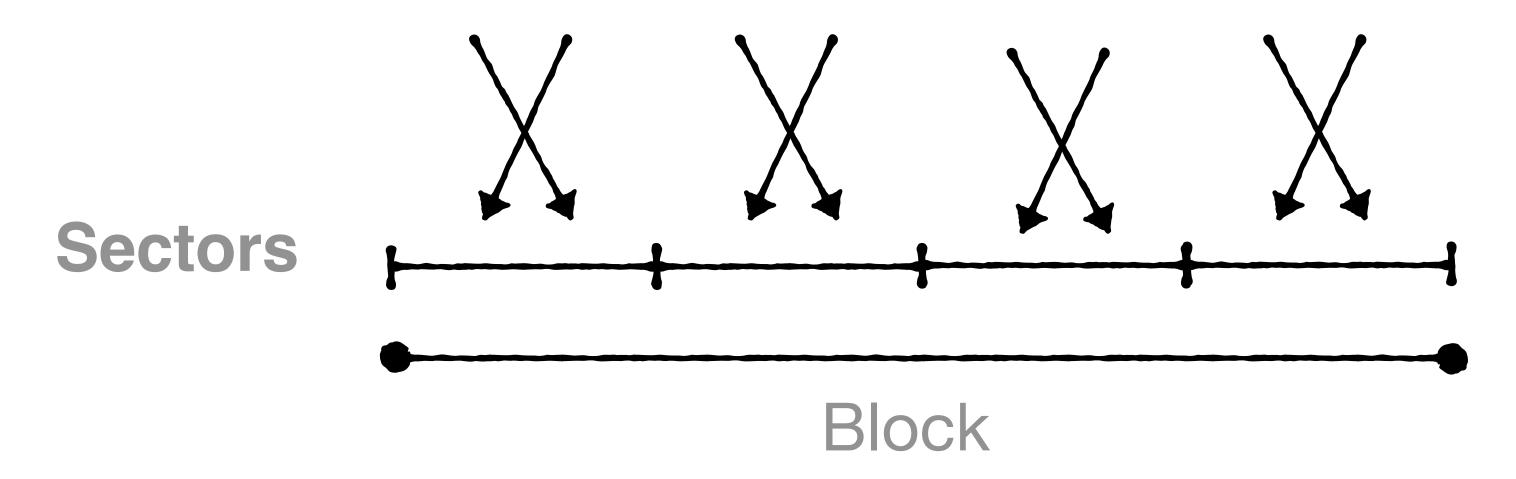
# Further ways to optimize? AVX - SIMD

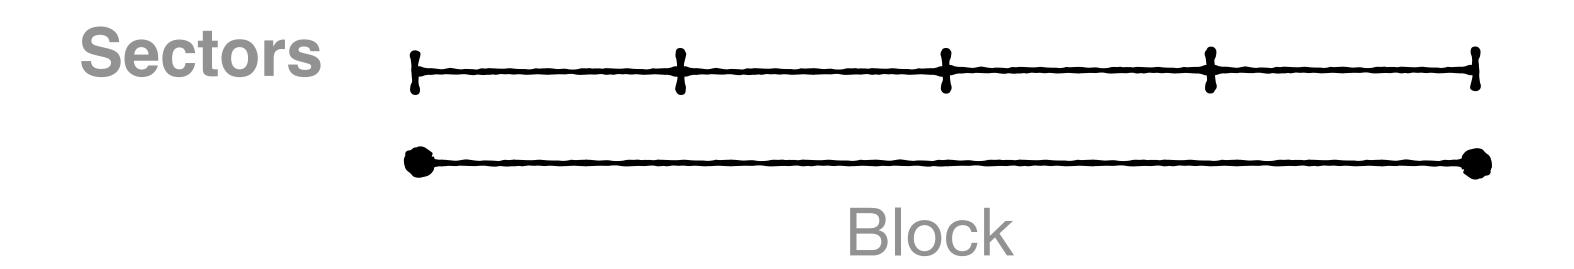


### Further ways to optimize?

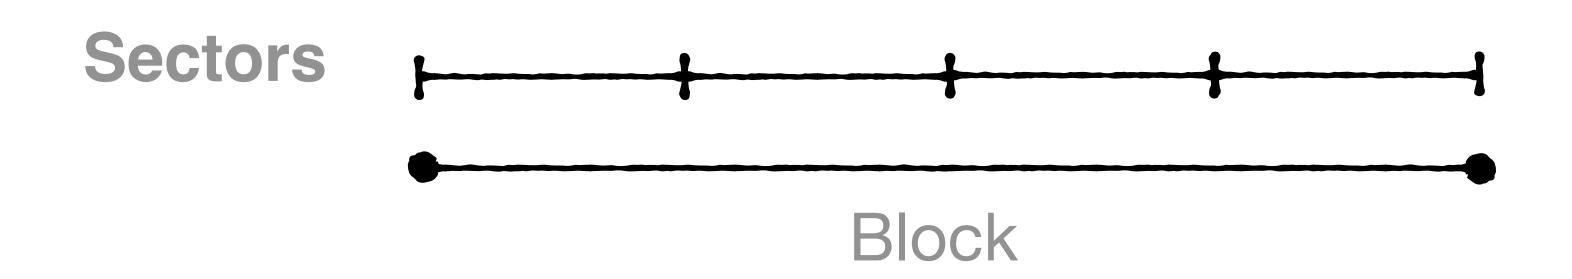
#### AVX - SIMD

### Operate on each sector in parallel



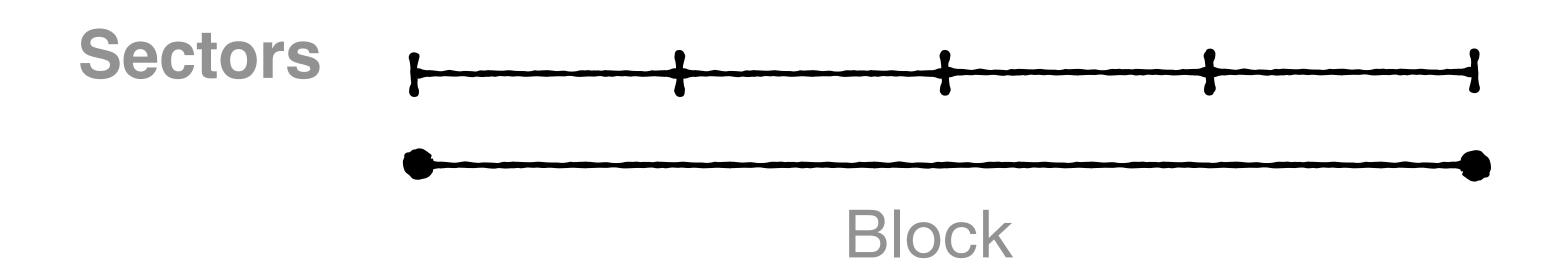


# # hashes must be multiple of # sectors



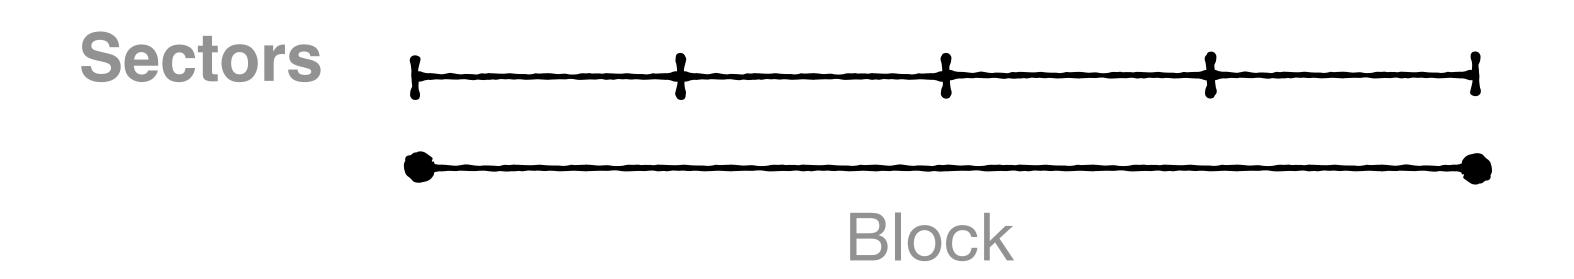
# hashes must be multiple of # sectors

Why is this bad?



# hashes must be multiple of # sectors Why is this bad?

Force using sub-optimal # hash functions

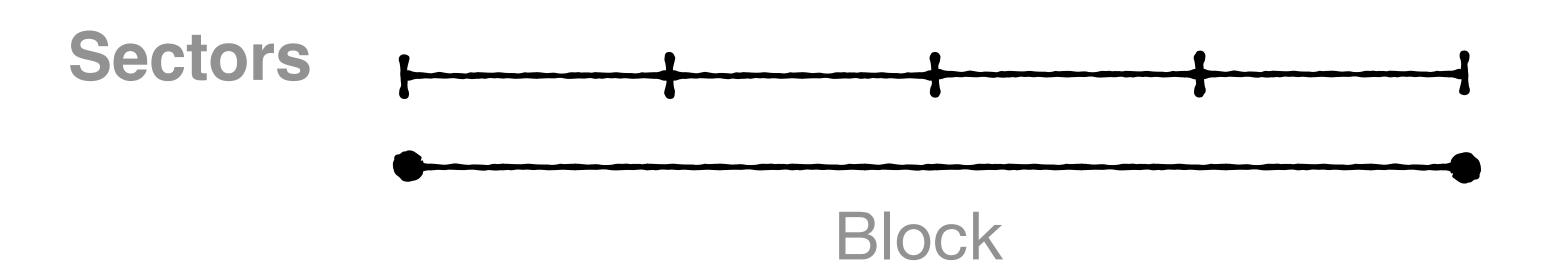


# hashes must be multiple of # sectors

Why is this bad?

Force using sub-optimal # hash functions

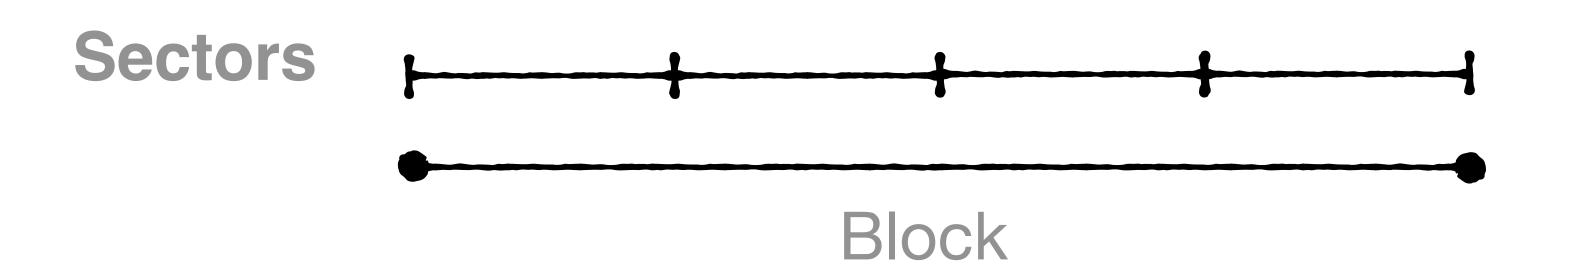
#### Harm FPR



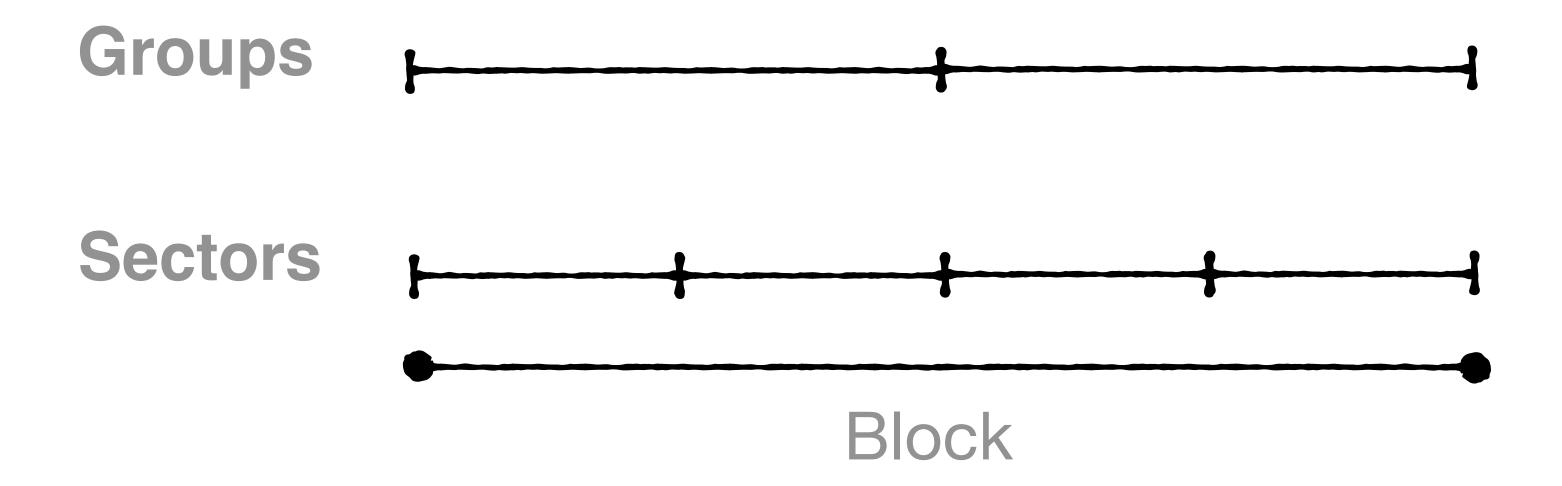
# hashes must be multiple of # sectors Why is this bad?

Force using sub-optimal # hash functions

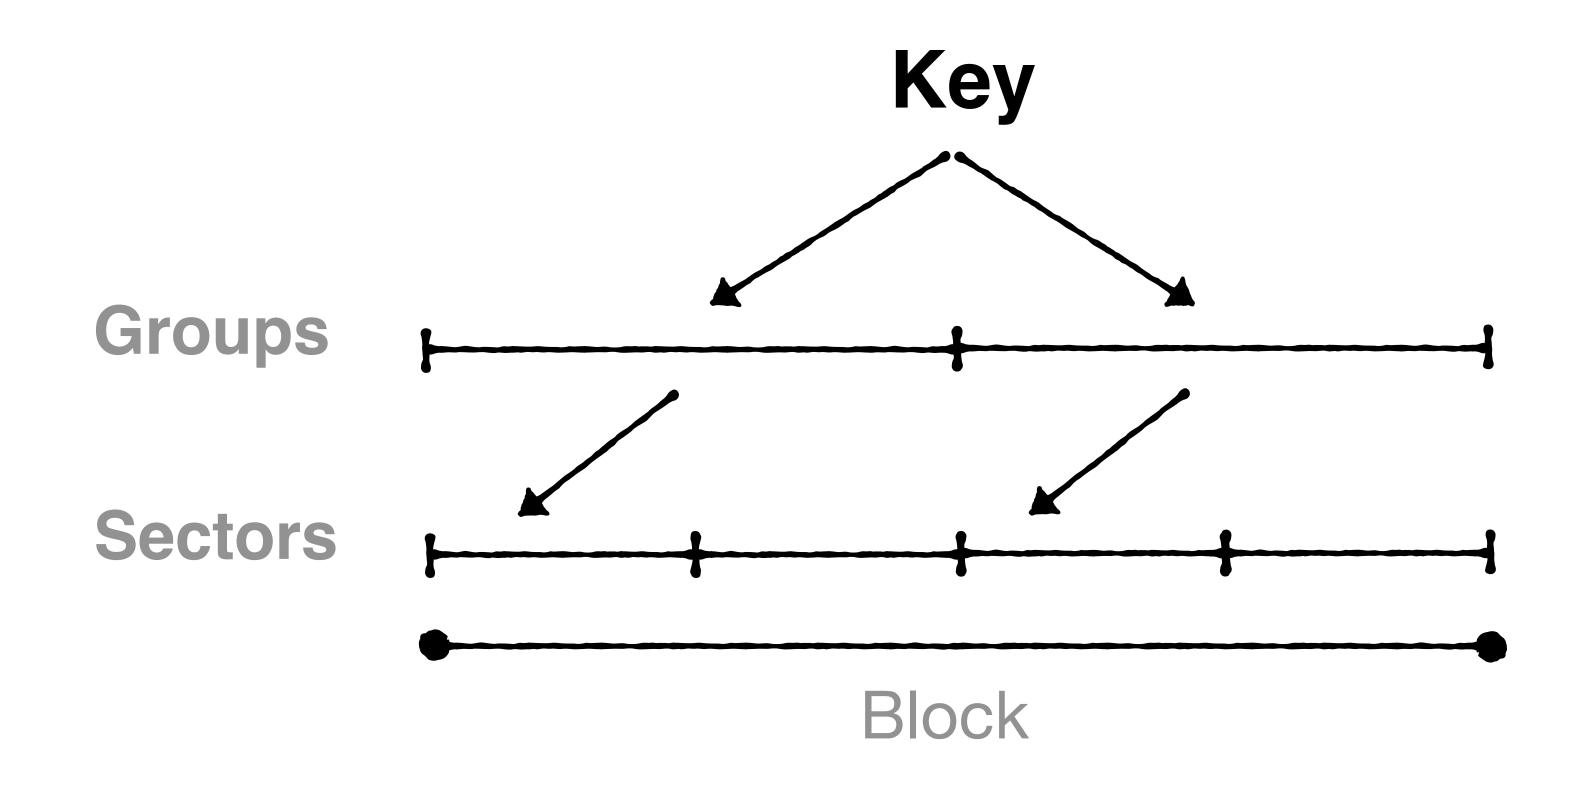
Harm FPR - solutions?



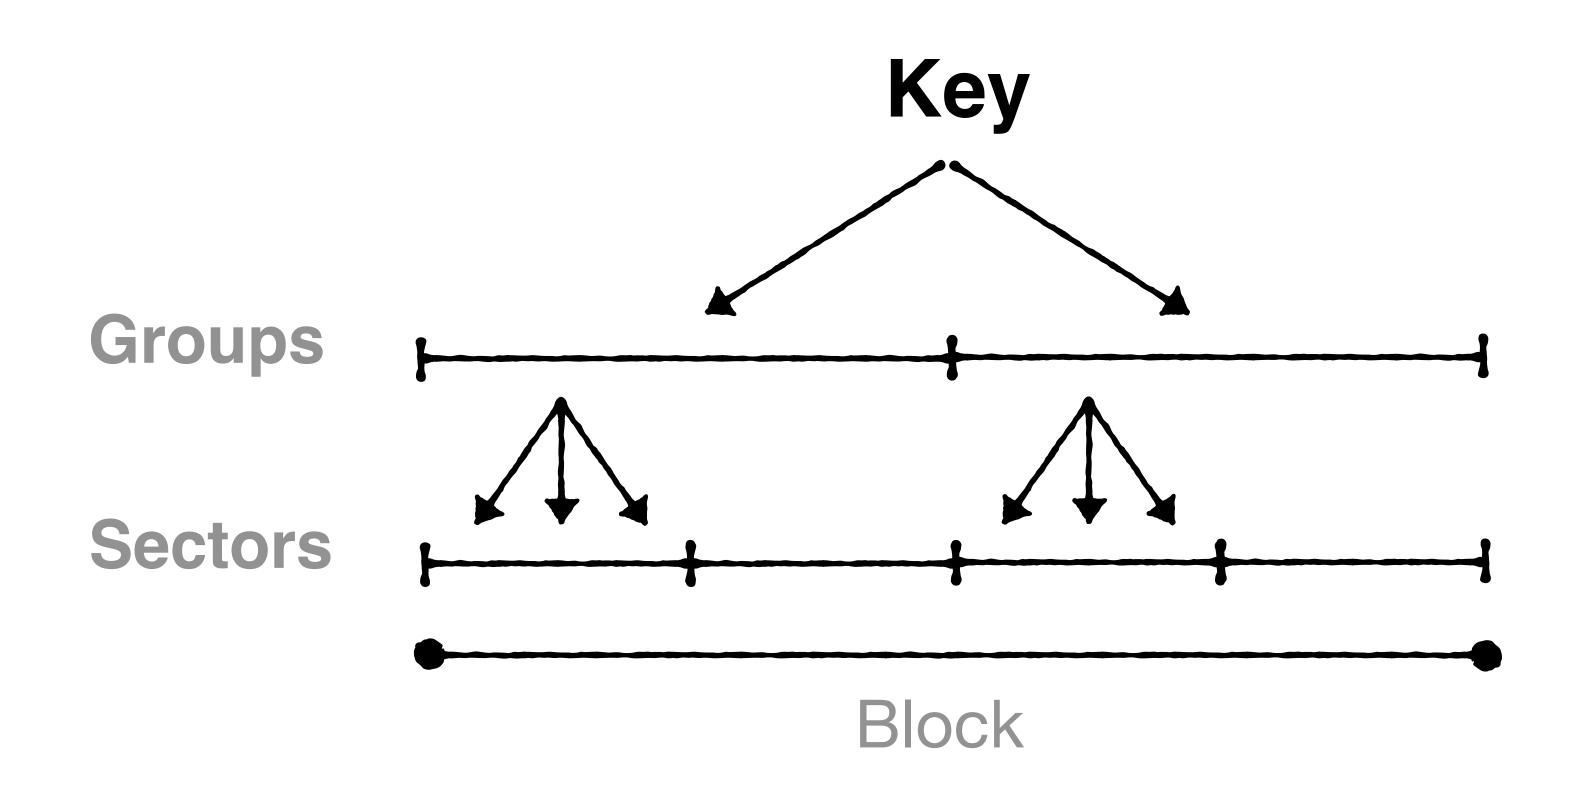
# Divide sectors into Z groups



# Map each key to one sector per group based on its hash

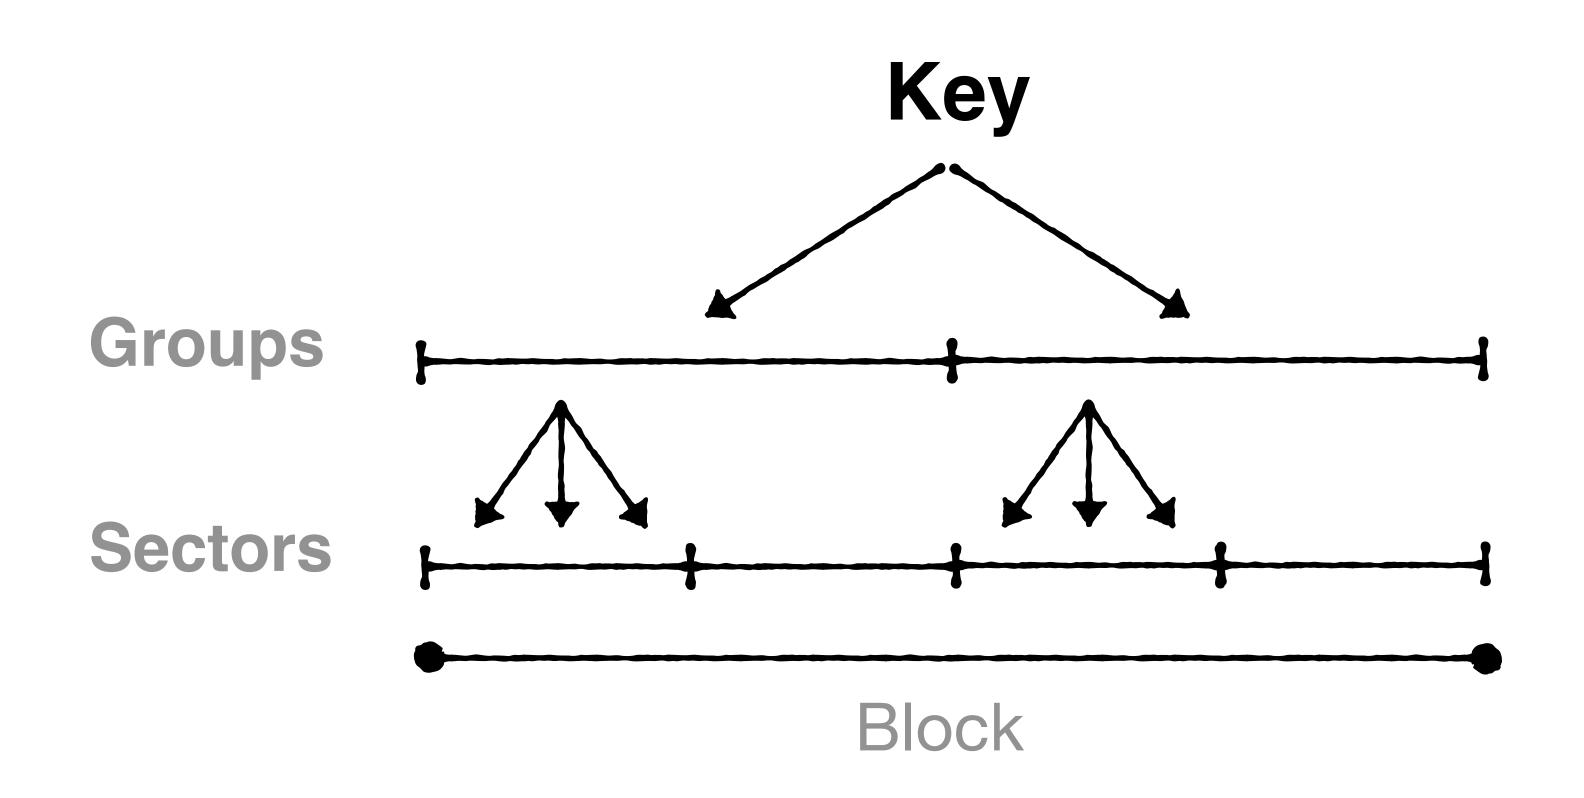


# Hash bits only to relevant sector in each group

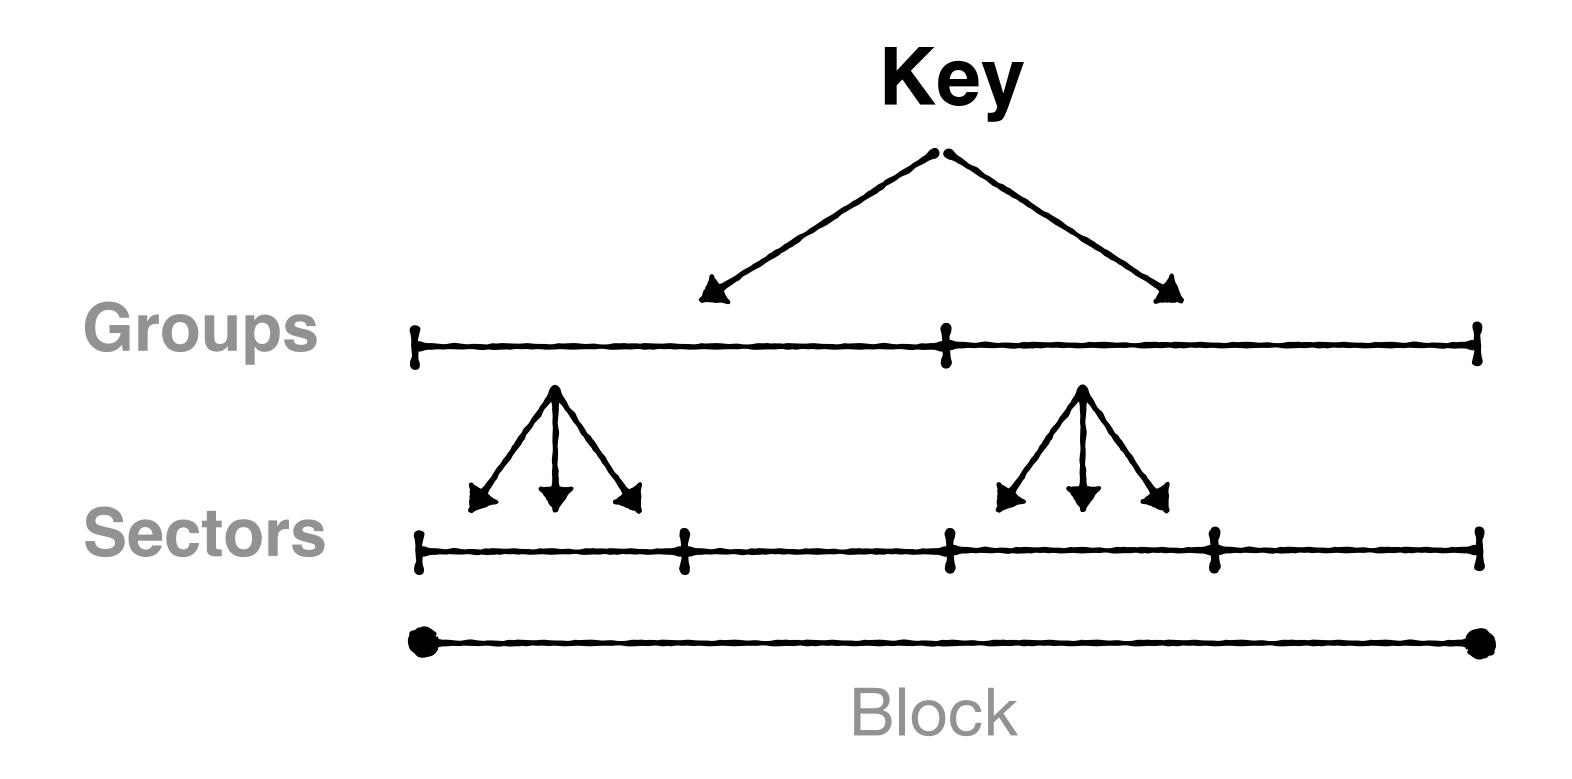


## Hash bits only to relevant sector in each group

# e.g., here we can use 6 hashes:)



Hash bits only to relevant sector in each group



Nearly best of both worlds:) fast & low FPR

# Break

## Bloom



FPR  $\varepsilon \approx 2 - M/N \cdot \ln(2)$ 

## Bloom



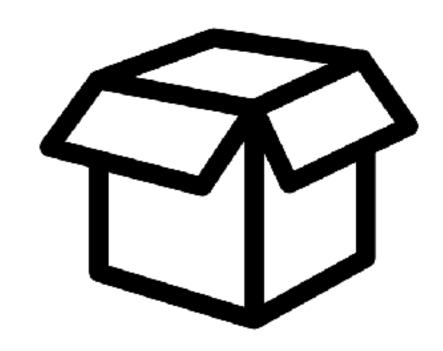
FPR  $\varepsilon \approx 2 - M/N \cdot \ln(2)$ 

#### Lower Bound

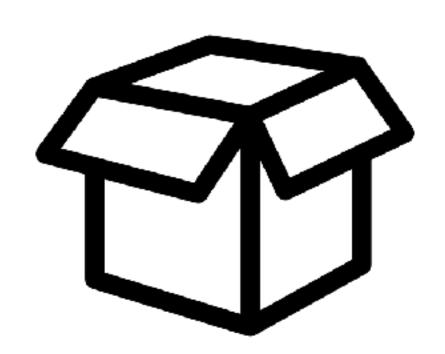


???

# Assume nothing about implementation



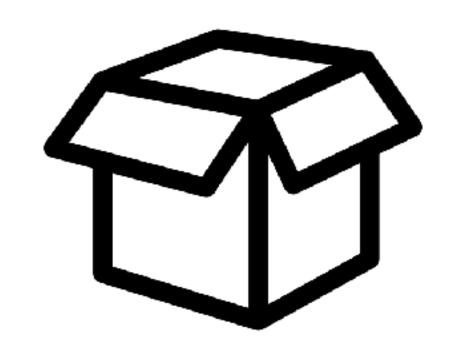
Assume nothing about implementation



# Analyze with respect to filter specification



Assume nothing about implementation



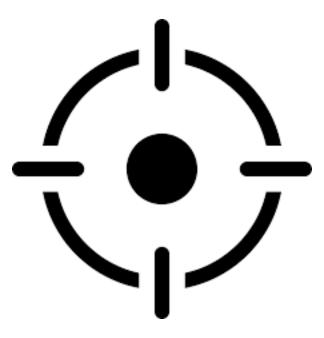
Analyze with respect to filter specification



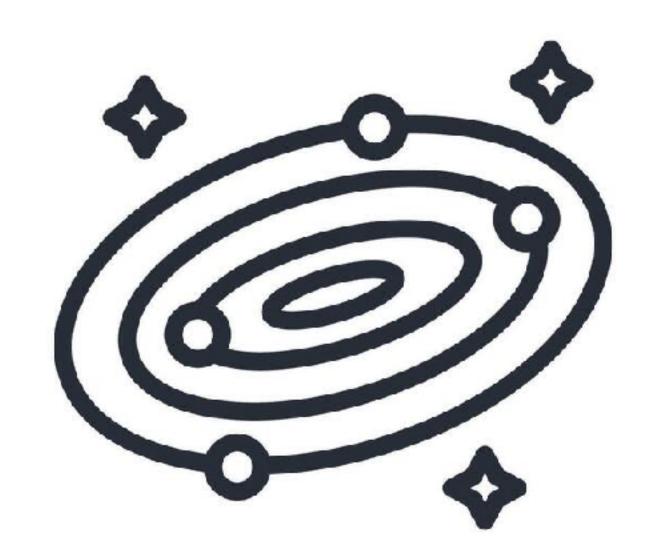
ε-FPR

N - # entries

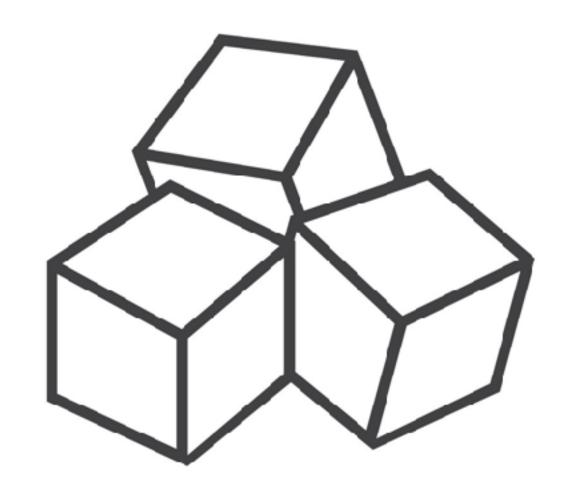
U - Universe size



# Out of Universe U, store N entries



# (U choose N) combinations



log<sub>2</sub>(U choose N) bits

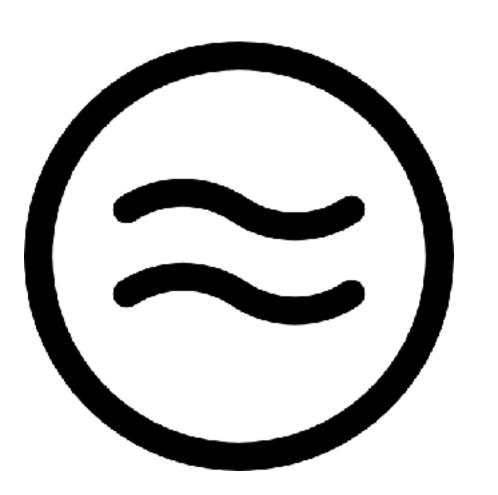
to encode a unique combination

 $log_2(U choose N) \approx N \cdot log_2(U/N)$  bits



for large U

# Lower Bound for Filter - Approximate Set



#### Lower Bound for Filter

IFilterl + IDisambiguationl ≥ IExact Setl

# IFilterl + IDisambiguationl ≥ IExact Setl

# Legend

ε-FPR N-#entries

# IFilterI + IDisambiguationI ≥ IExact SetI



# What information must we add the filter to turn it into an exact set?

# Legend

ε - FPR

N - # entries



Plug in

# Legend

ε - FPR

N - # entries



# Query filter U times

# Legend

ε - FPR

N - # entries



Tells us all positive keys

$$N + \epsilon \cdot U$$

# Legend

ε - FPR

N - # entries

1

Tells us all positive keys

 $\epsilon \cdot U$ 

# Legend

ε - FPR

N - # entries



# Of all positives ε · U, which keys are positive N

# Legend

ε - FPR

N - # entries



ε· U choose N

# Legend

ε - FPR

N - # entries



 $log_2(\varepsilon \cdot U choose N)$ 

# Legend

ε-FPR N-#entries



 $N \cdot log_2((\epsilon \cdot U) / N)$ 

# Legend

ε - FPR

N - # entries

IFilter I +  $N \cdot log_2((\epsilon \cdot U) / N) \ge N \cdot log_2(U / N)$ 

# Legend

ε - FPR

N - # entries

IFilter  $\geq N \cdot \log_2(U/N) - N \cdot \log_2((\epsilon \cdot U)/N)$ 

# Legend

ε - FPR

N - # entries

# IFilter $\geq N \cdot \log_2(1/\epsilon)$



# Legend

ε - FPR

N - # entries

 $\epsilon \geq 2 - M/N$ 



# Legend

ε-FPR

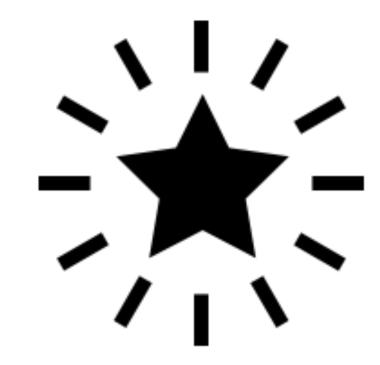
N - # entries

### Bloom



 $\approx 2 - M/N \cdot 0.69$ 

#### Lower bound



2 -M/N

## Bloom



 $\approx 2 - M/N \cdot 0.69$ 

#### Lower bound





2 -M/N

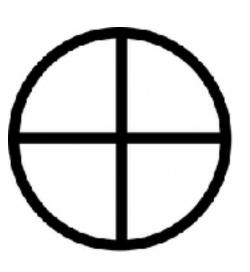
Bloom

**XOR Filter** 

Lower bound





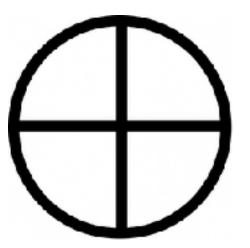


≈ 2 -M/N · 0.81



 $\approx 2 - M/N$ 

#### **XOR** Filter

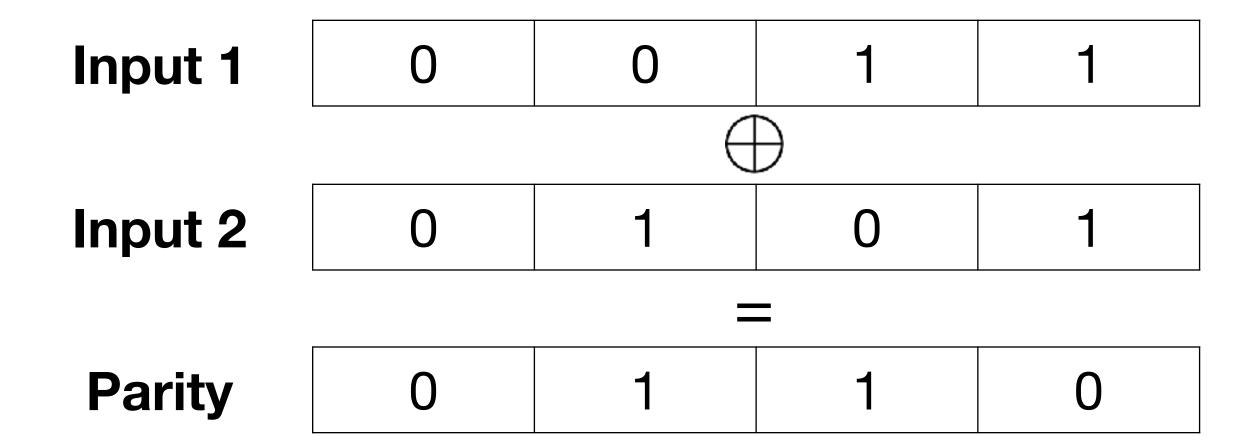


#### Xor Filters: Faster and Smaller Than Bloom Filters

Thomas Mueller Graf, Daniel Lemire

Journal of Experimental Algorithmics, 2020

## **XOR Operator**



Suppose we lost input 2

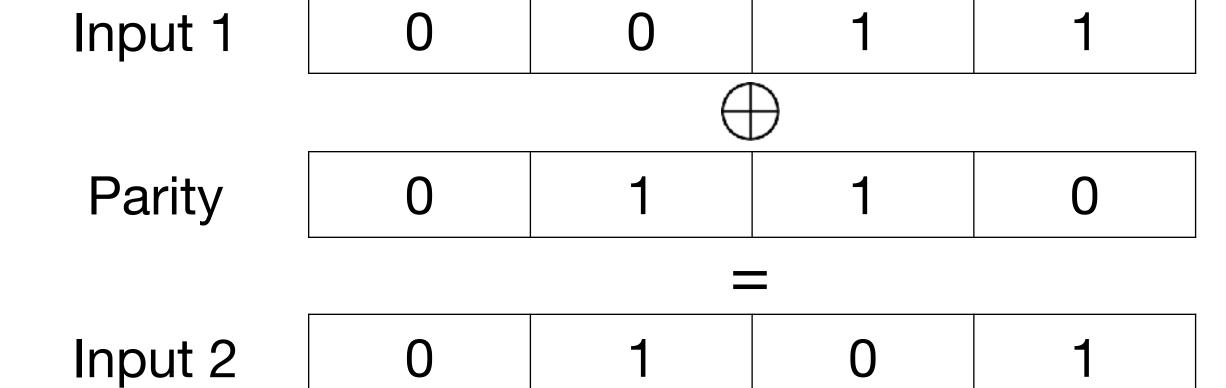
Input 1

0 0 1 1

Parity

0 1 1 0





Recovered

Or suppose we lost input 1

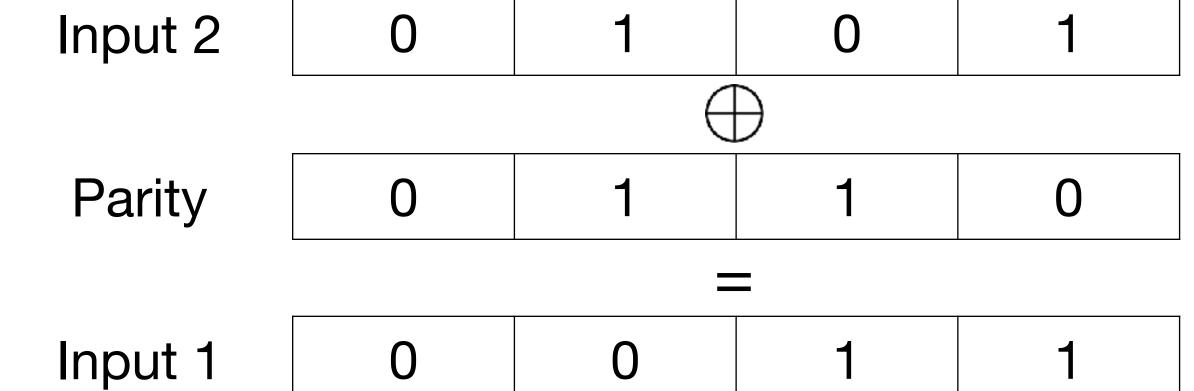
Input 2

0 1 0 1

Parity

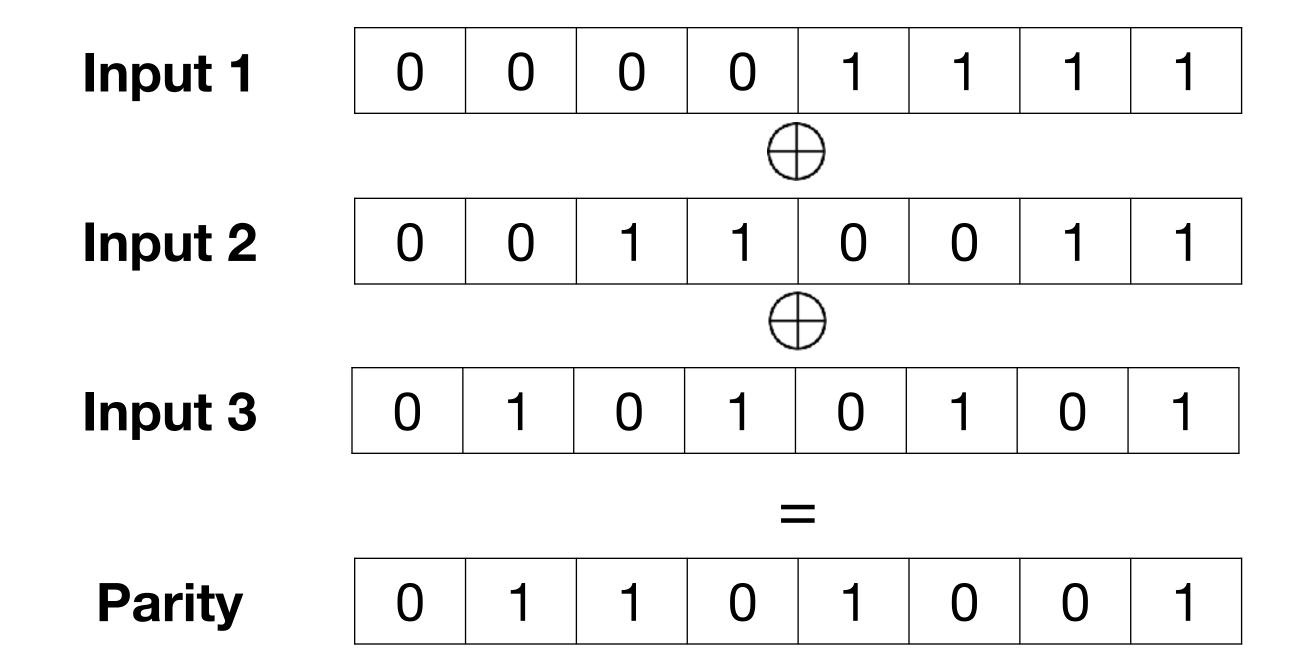
0 1 1 0



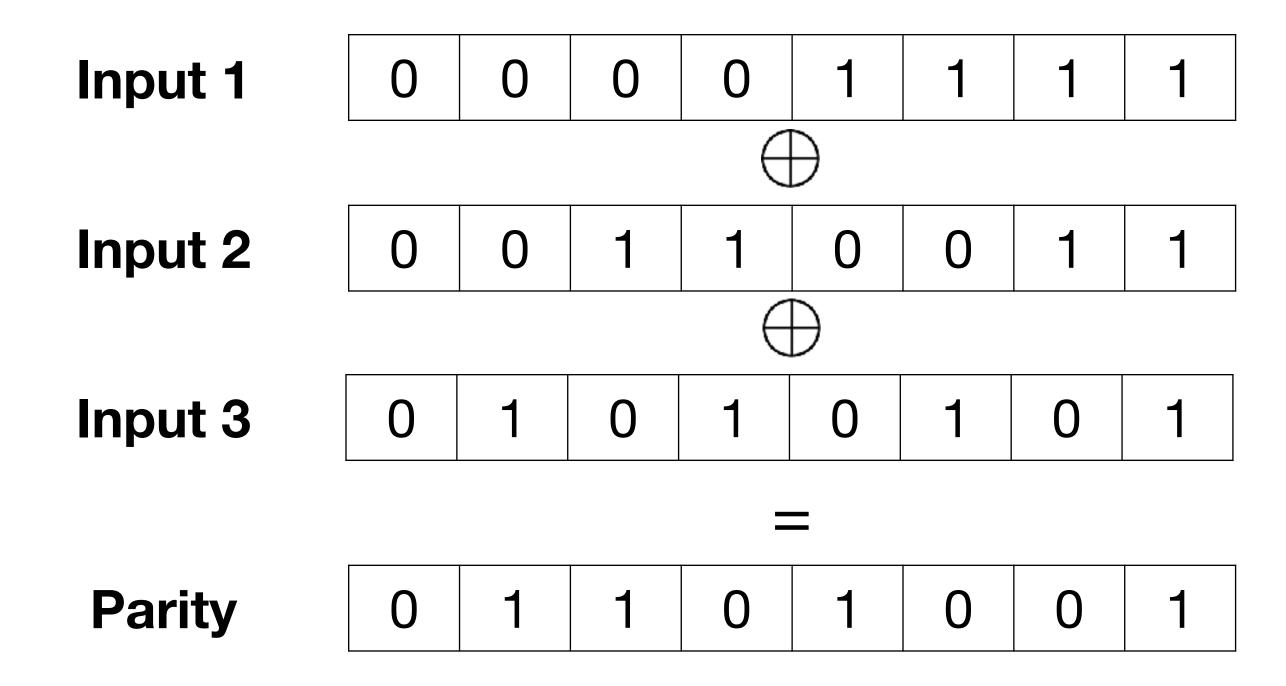


Recovered

#### XOR is commutative and associative



#### XOR is commutative and associative

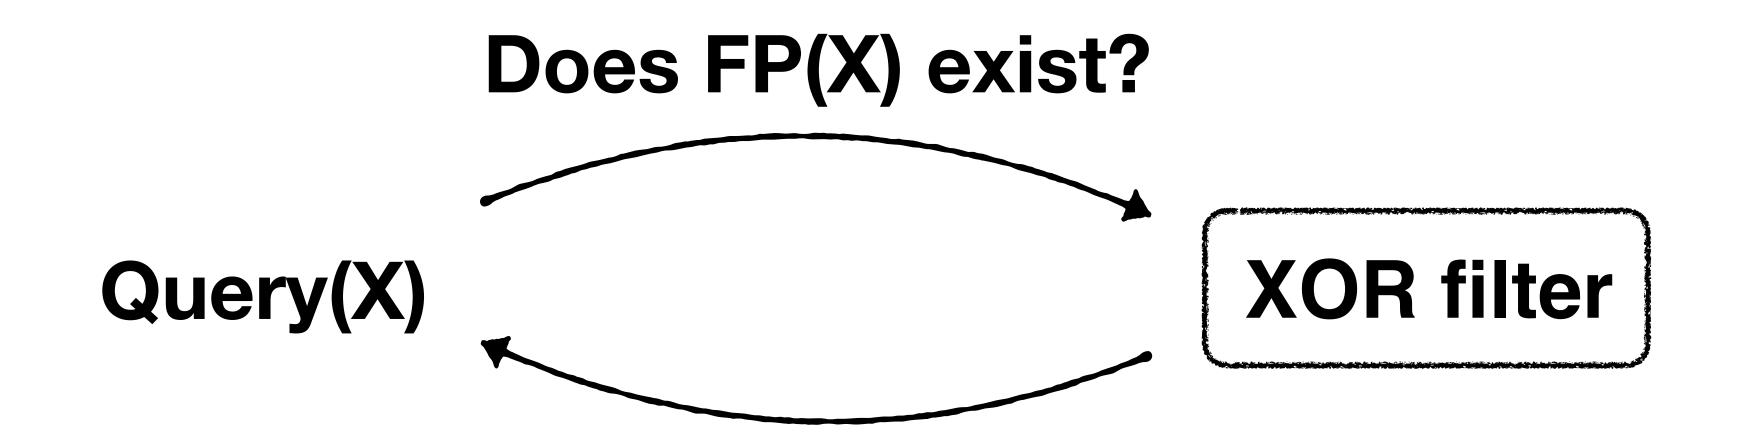


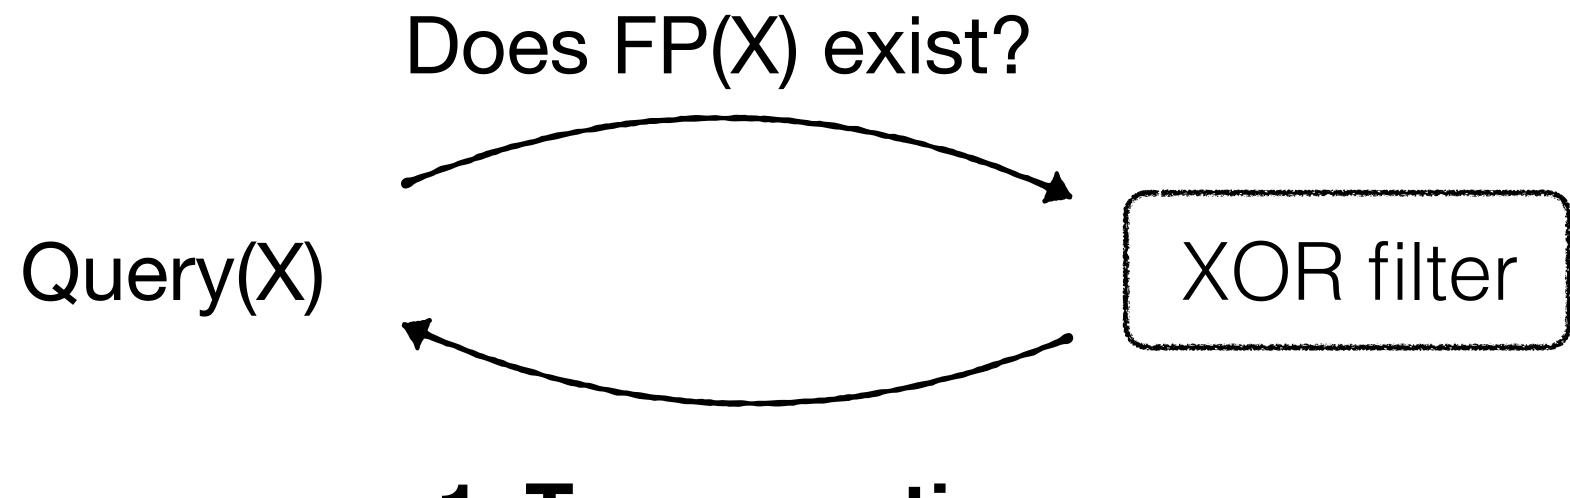
Parity can recover any input, as long as we also have all the other inputs

# XOR filter stores a fingerprint for each key

# XOR filter stores a fingerprint for each key

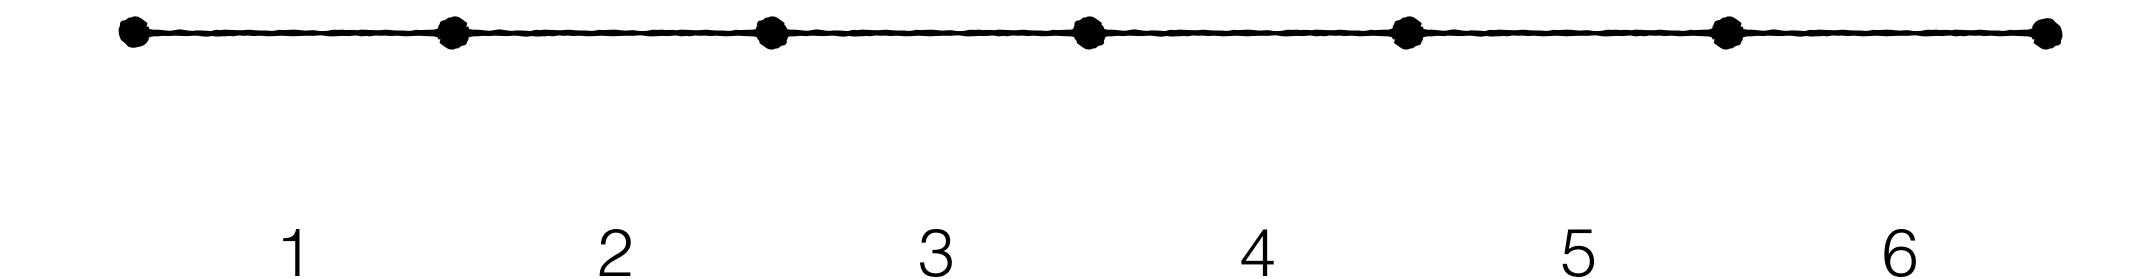
Example: 
$$FD(X) = 0100$$



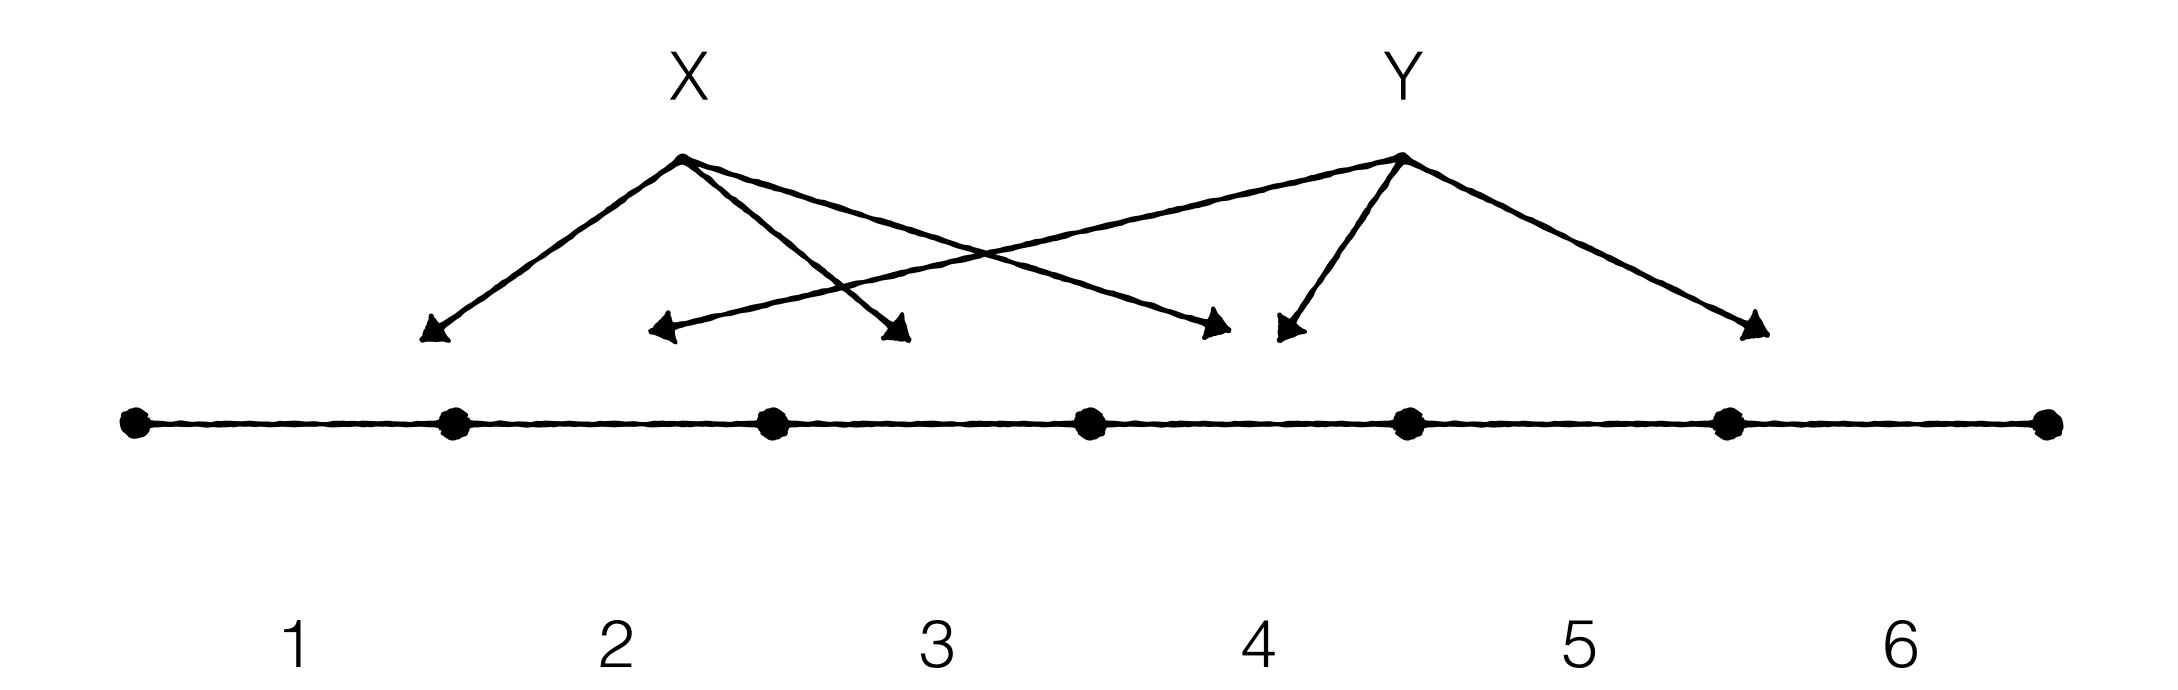


- 1. True negative
- 2. True positive
- 3. False positive with probability 2-F

# How does XOR filter store its fingerprints?

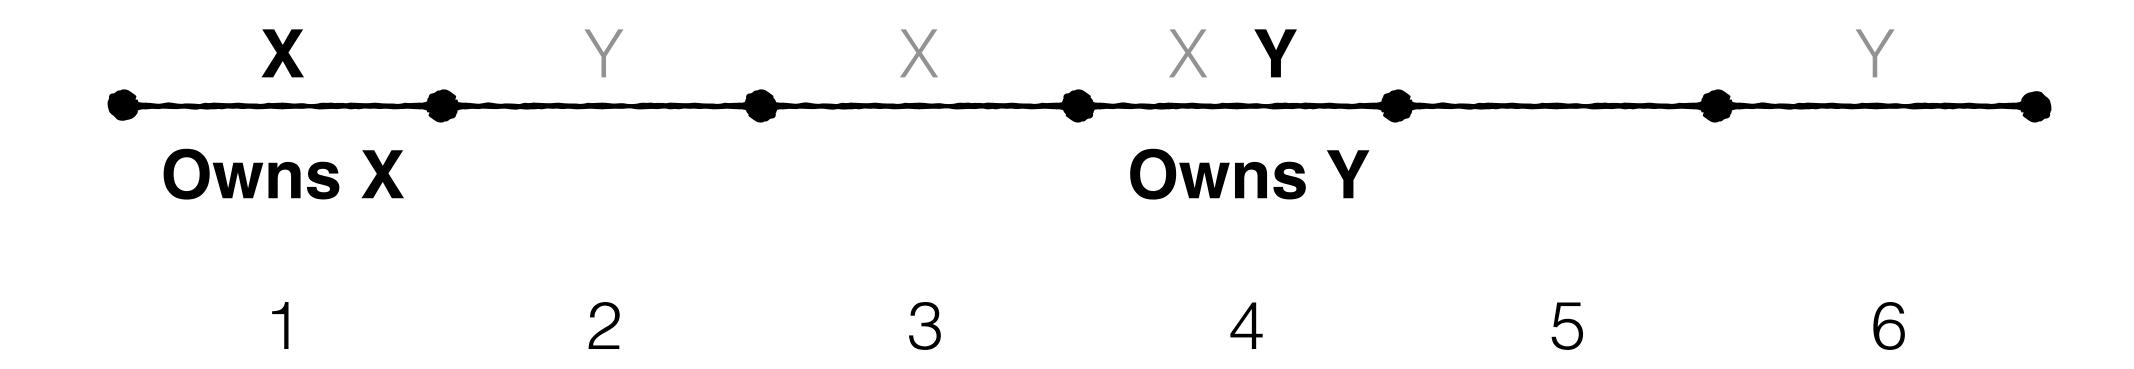


# Hash each entry to three slots

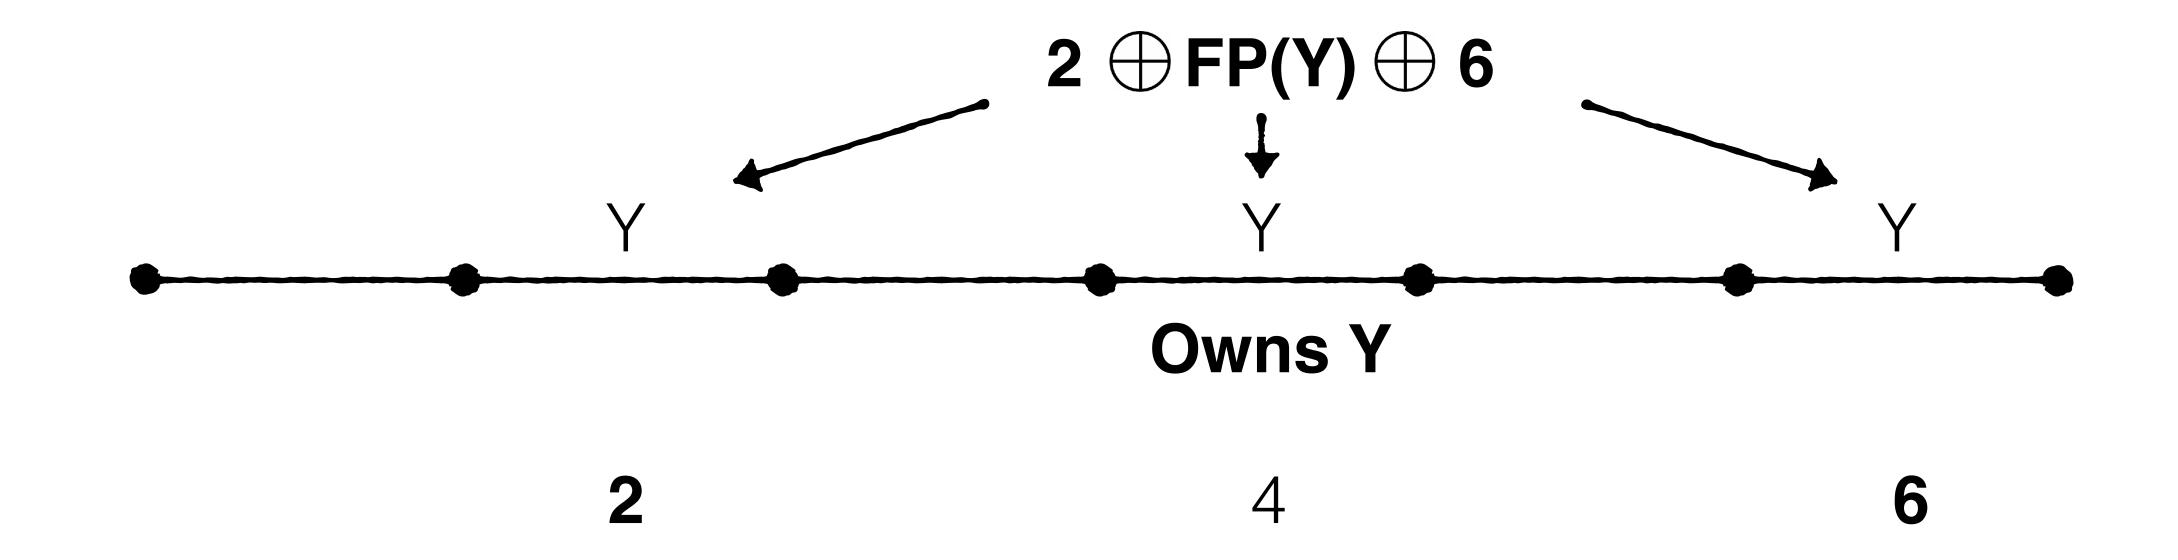


# Assign one slot to uniquely own each entry

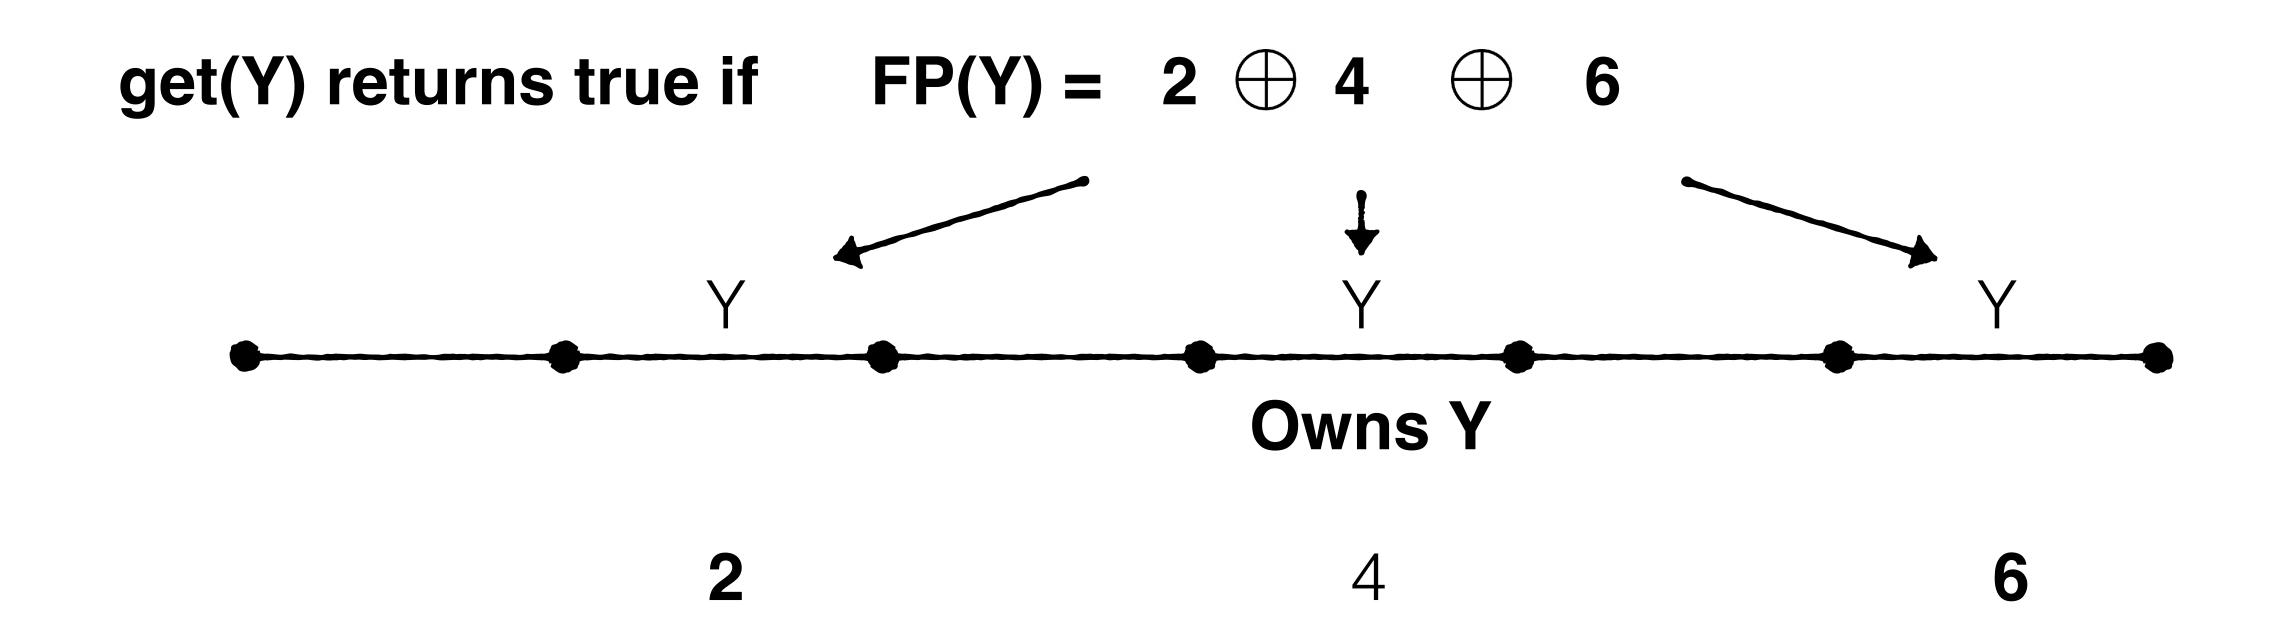
(More on this shortly)



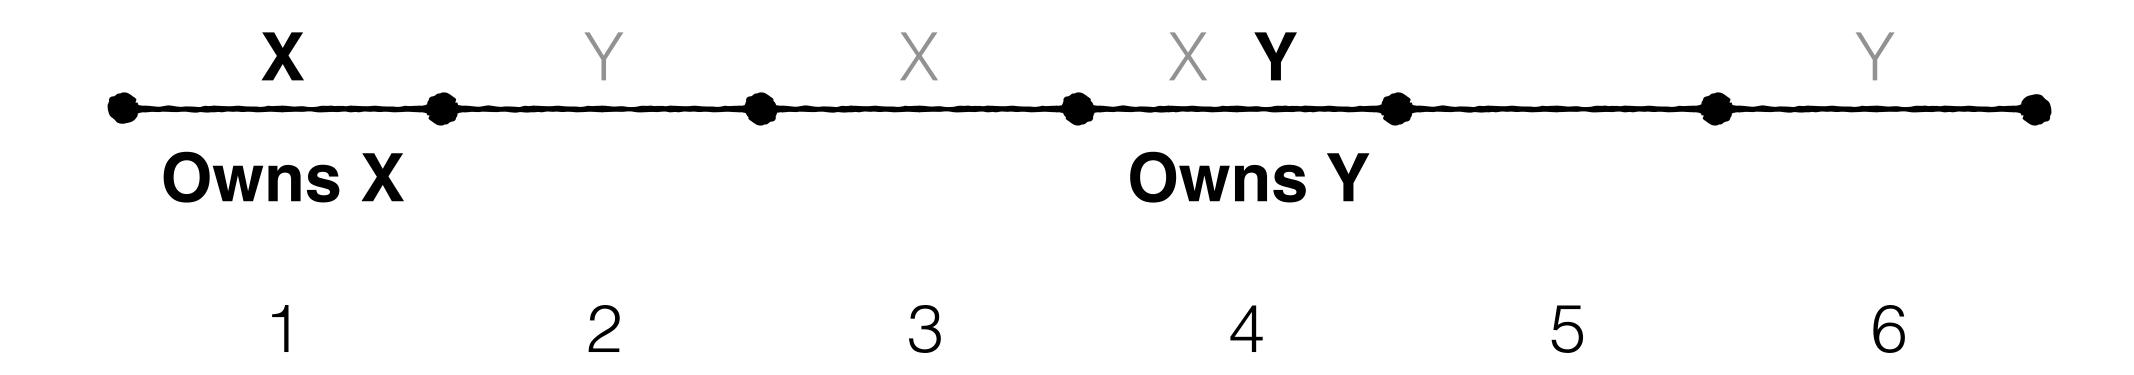
# Each bucket stores XOR of fingerprint and other two slots



# During queries, recover fingerprints by xoring three slots



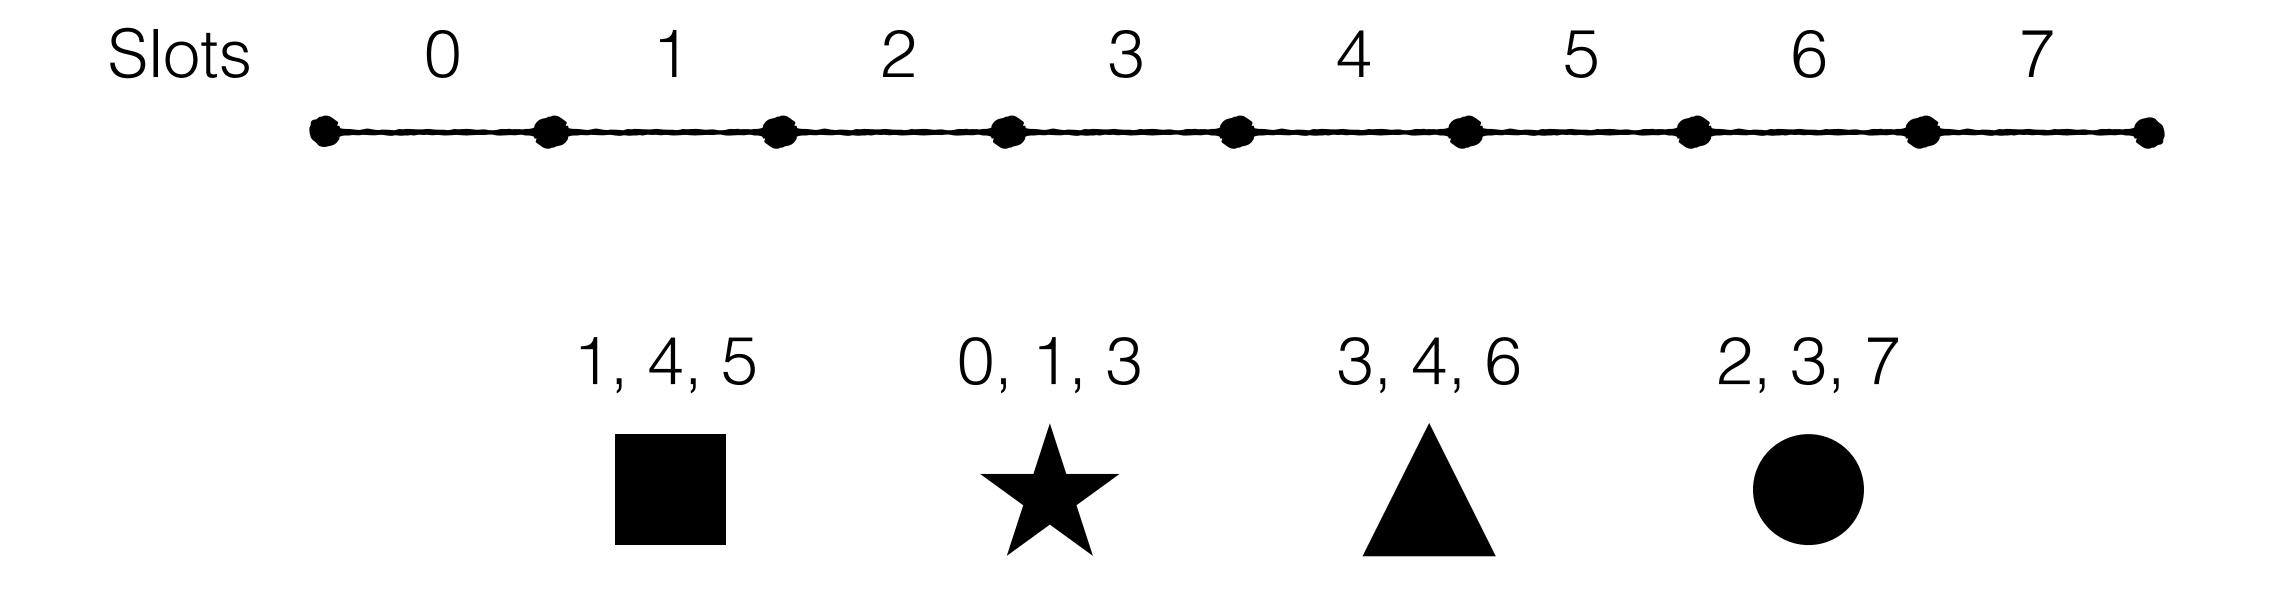
## How to assign slots to own different entries?

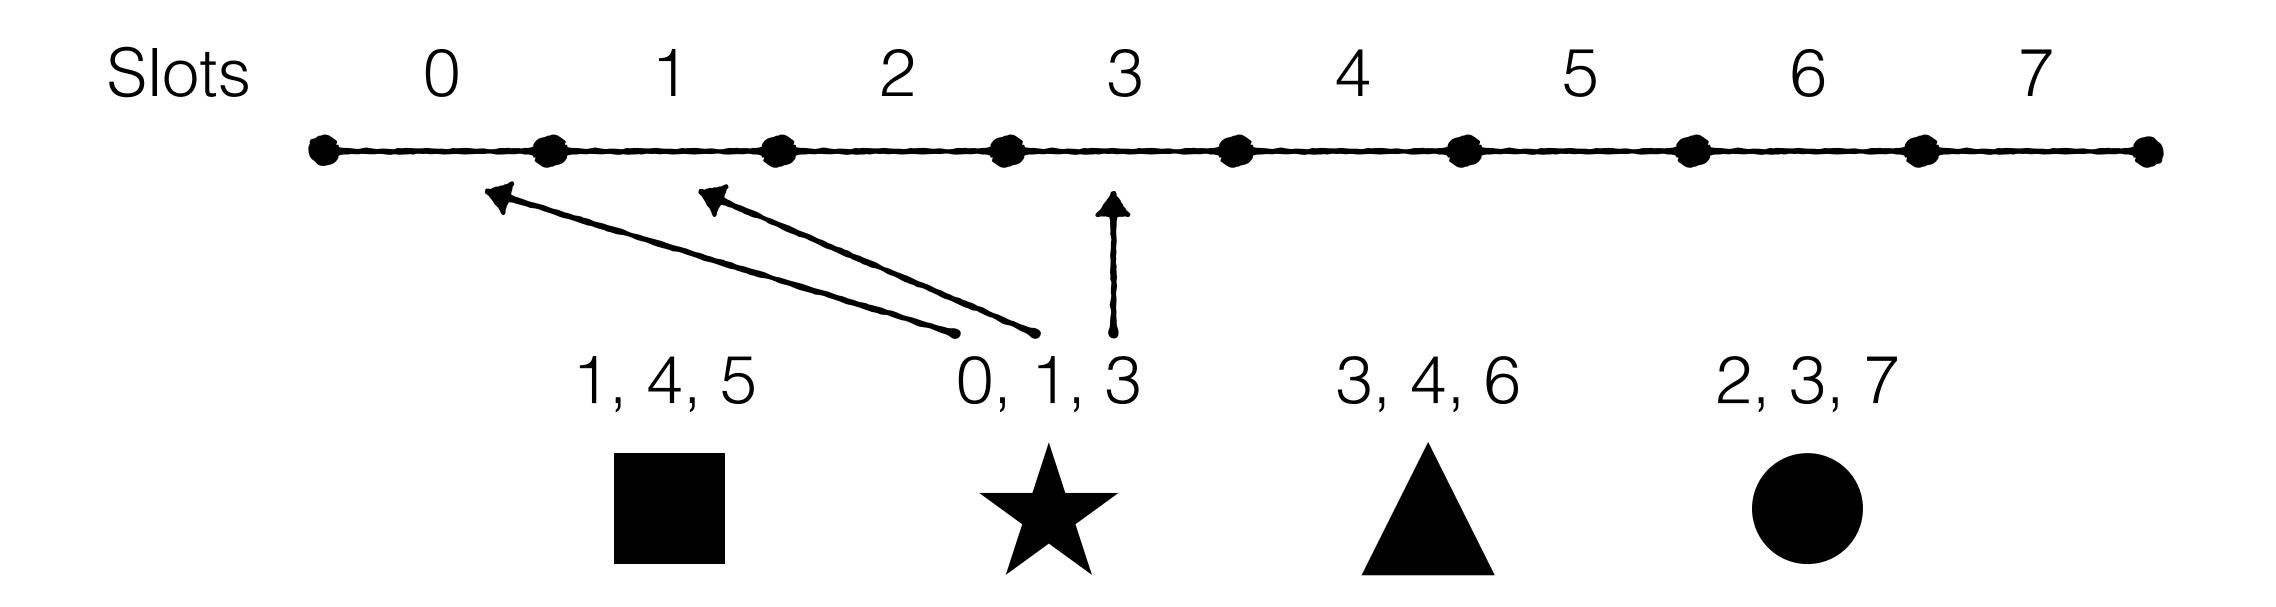


How to assign slots to own different entries?

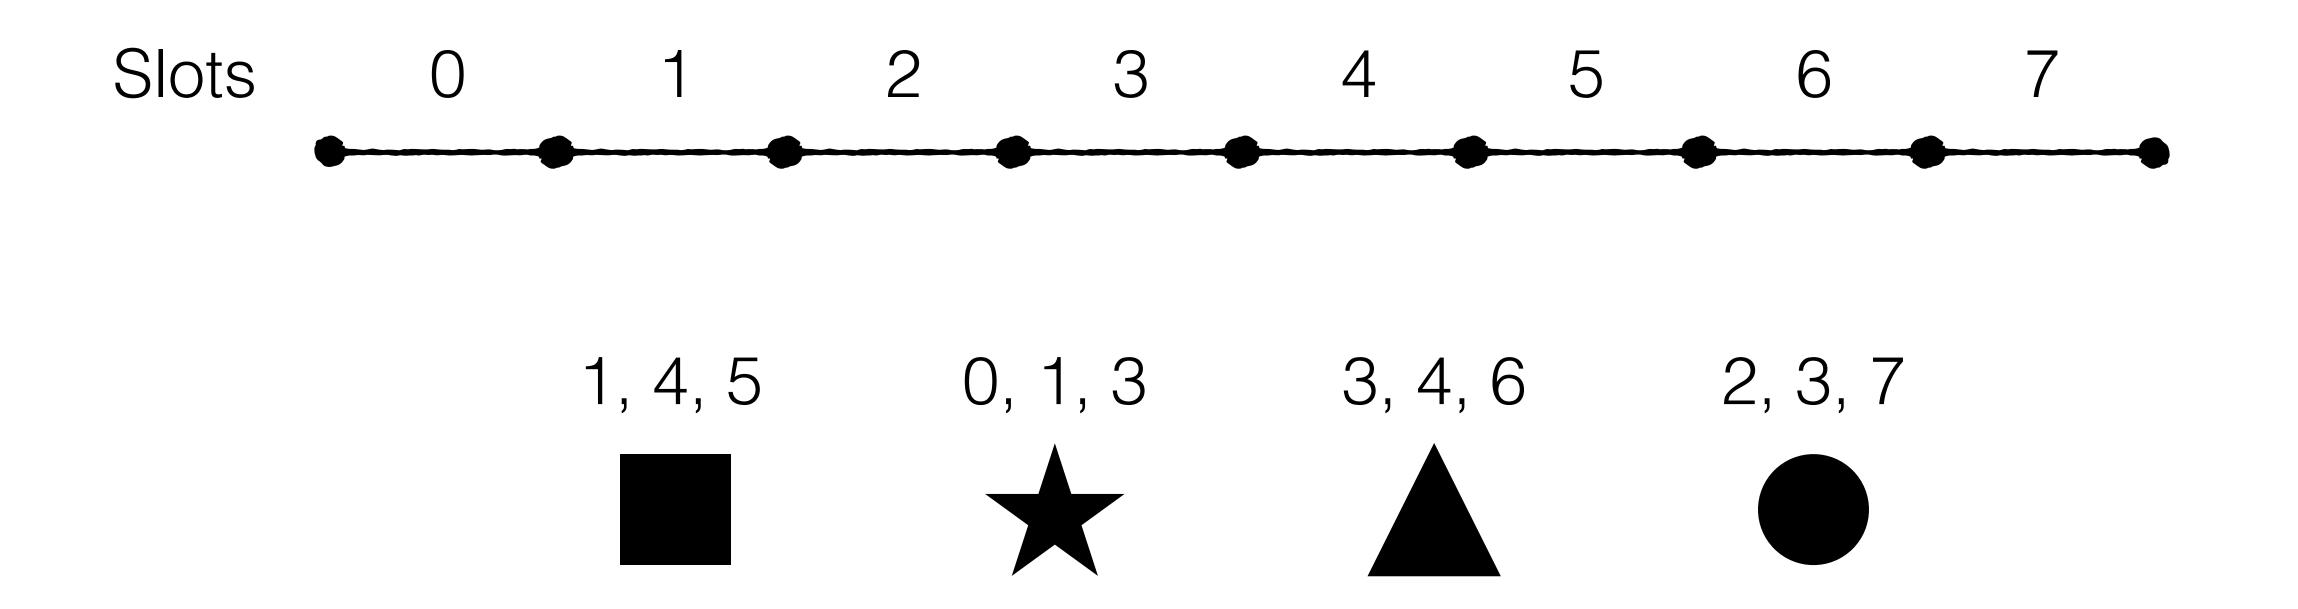


Peeling

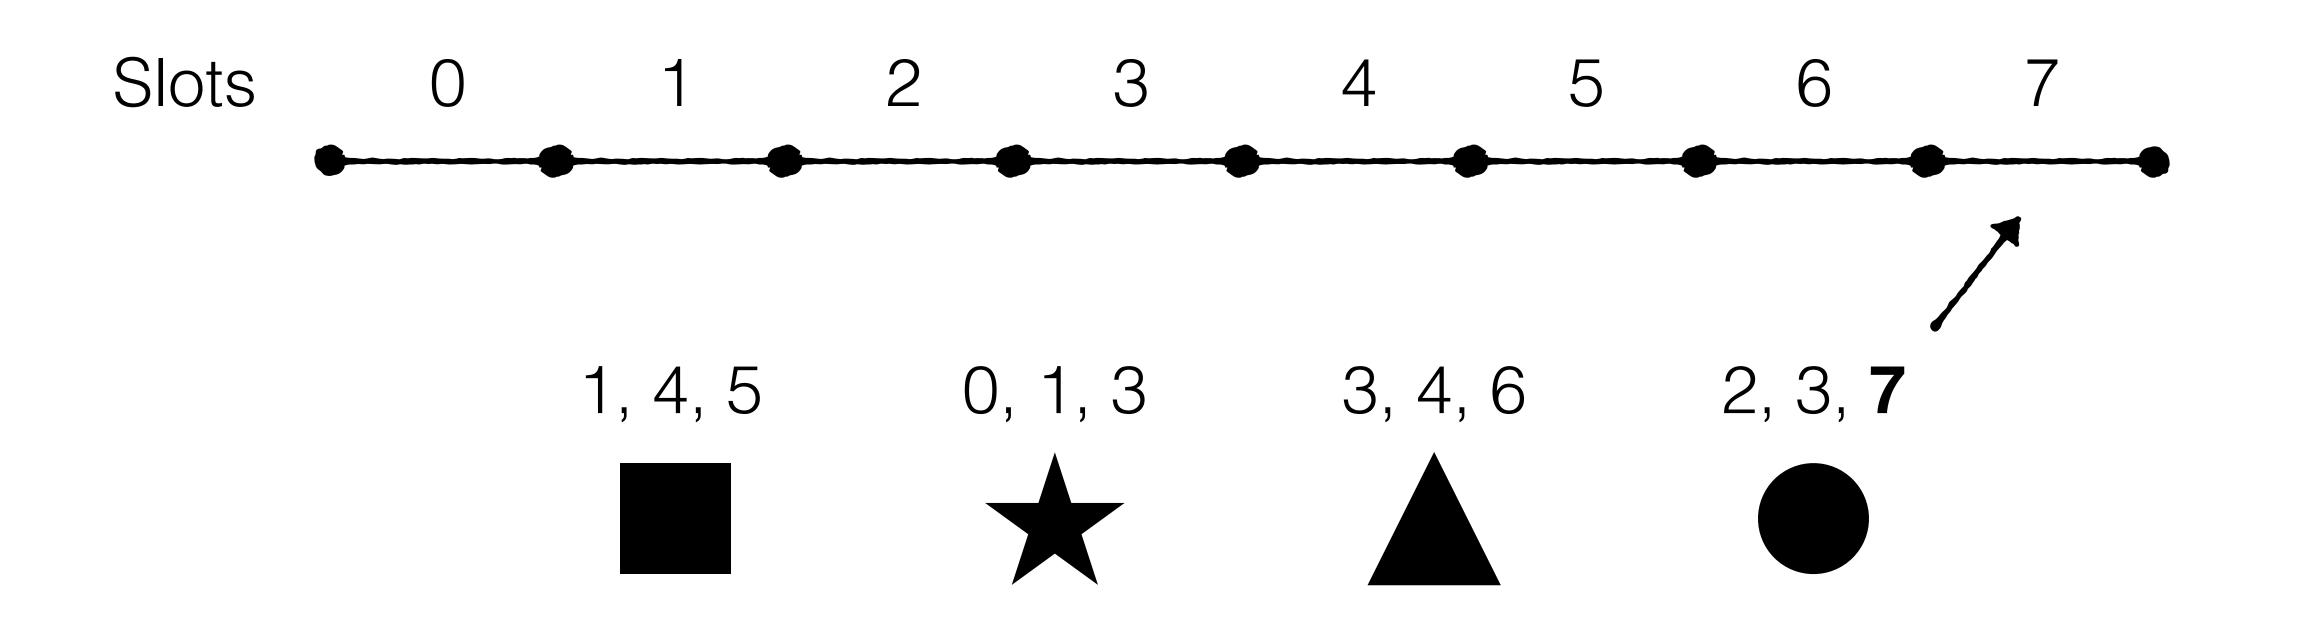




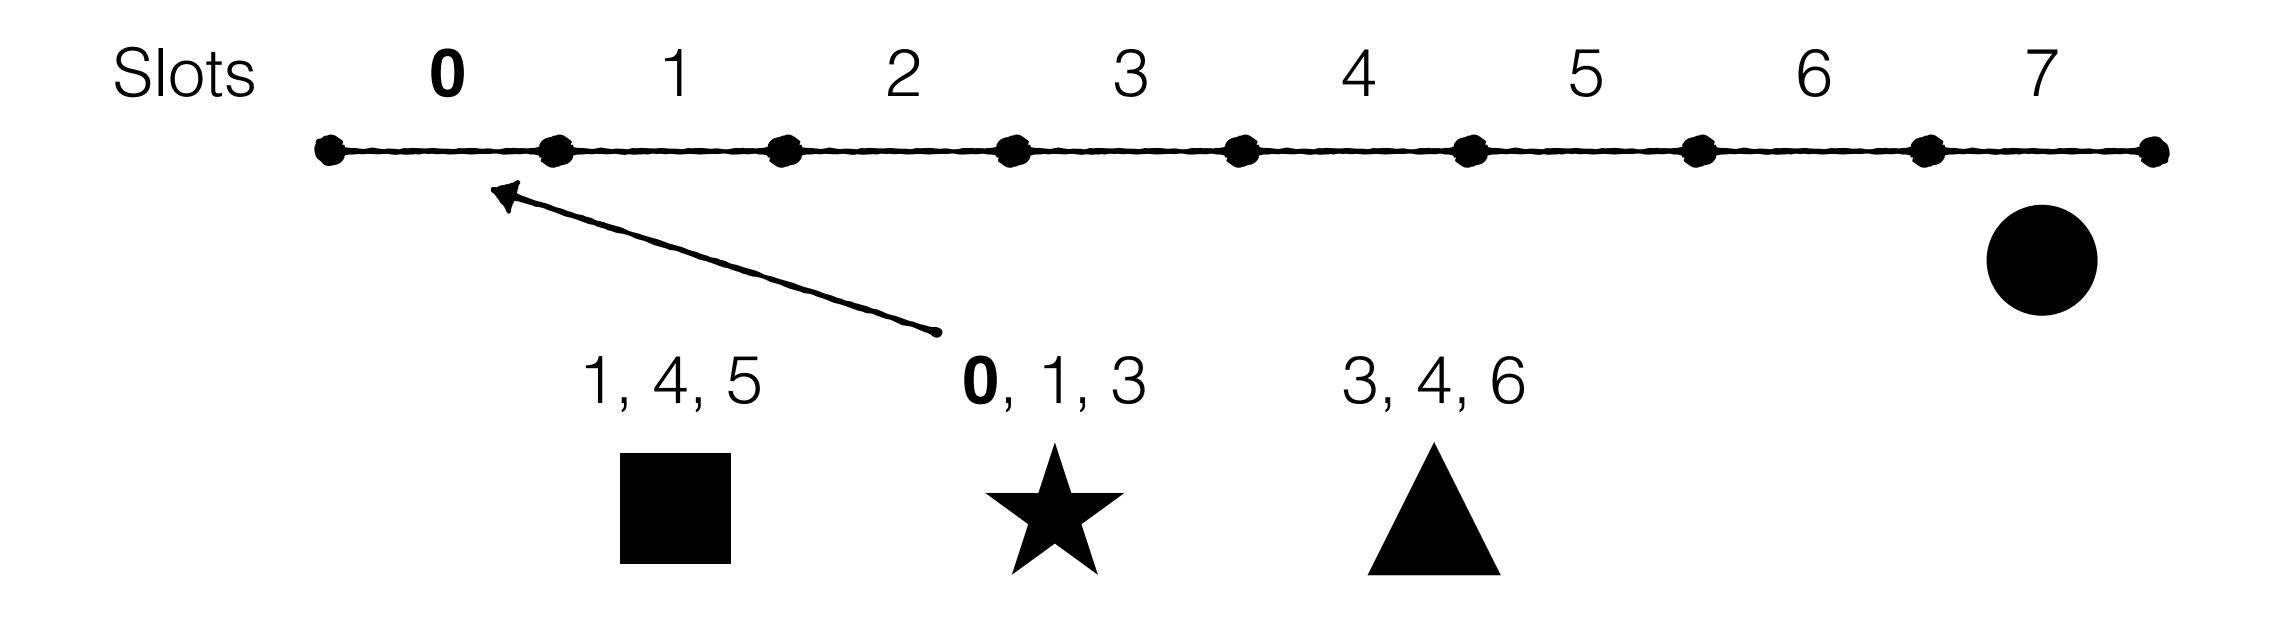
#### While not all keys have been assigned to a slot



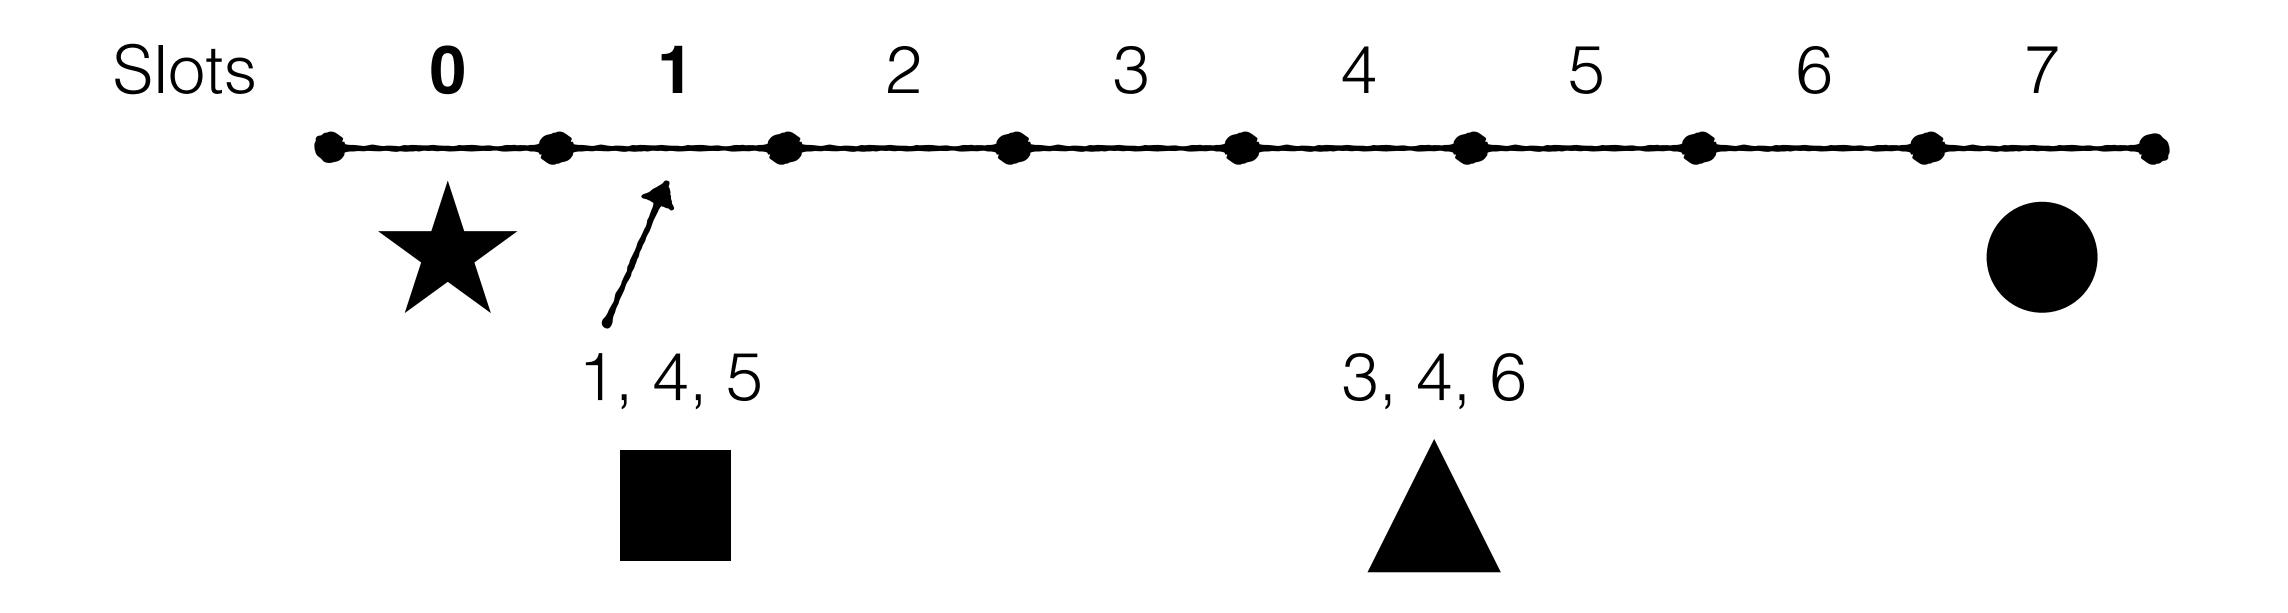
While not all keys have been assigned to a slot



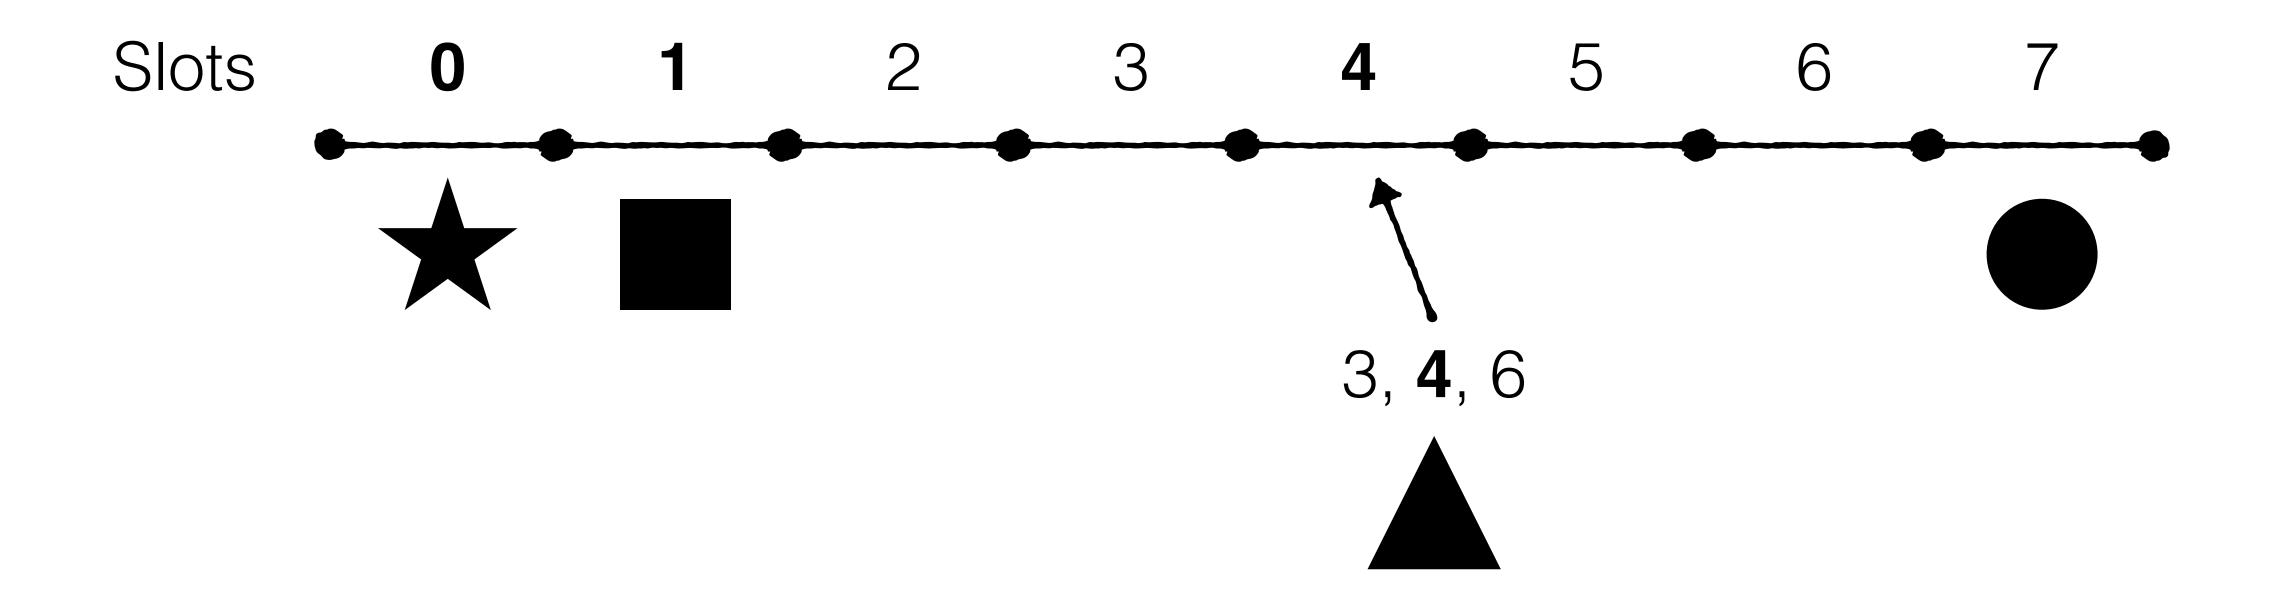
While not all keys have been assigned to a slot

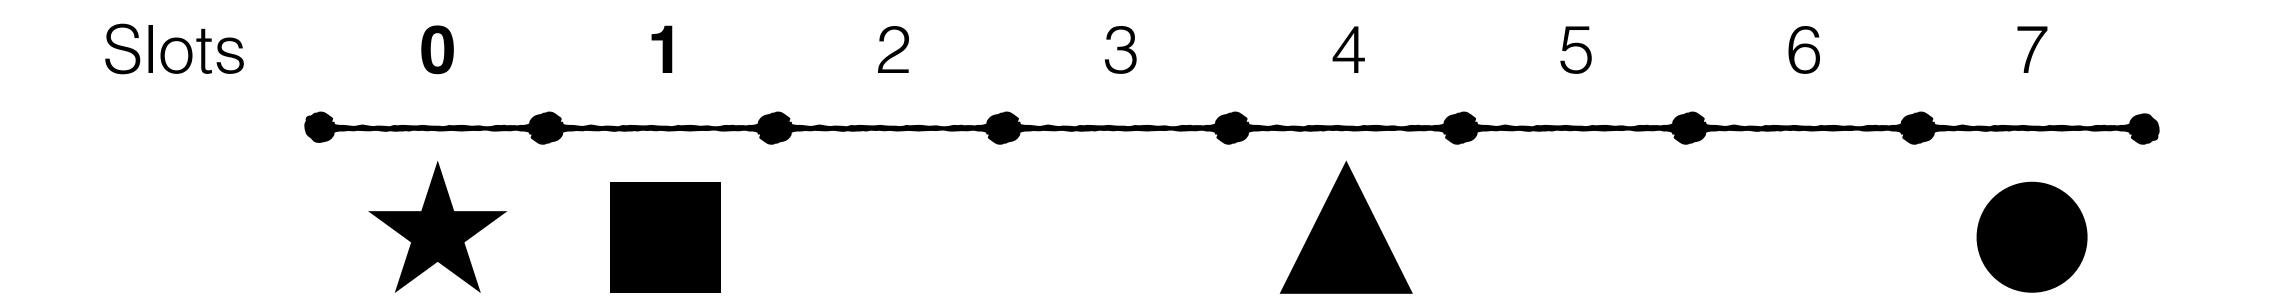


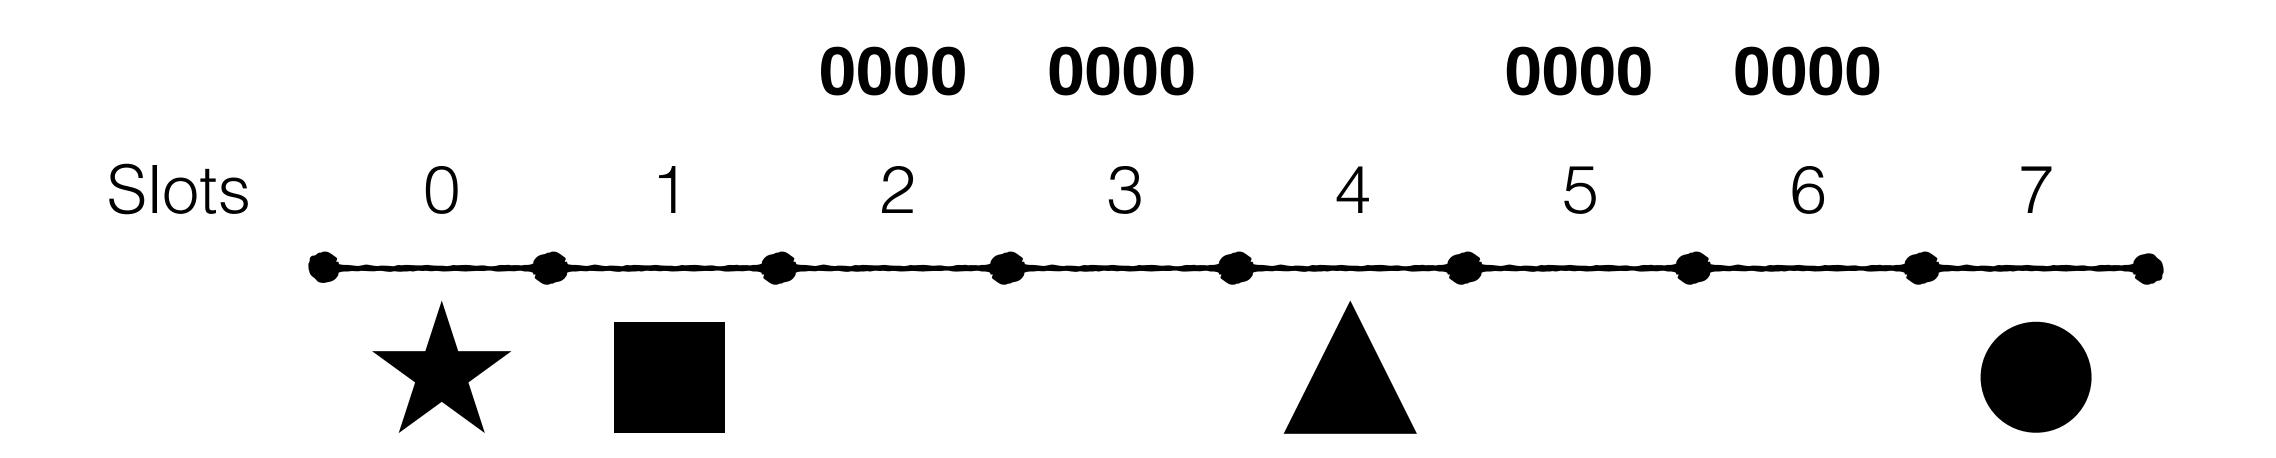
While not all keys have been assigned to a slot



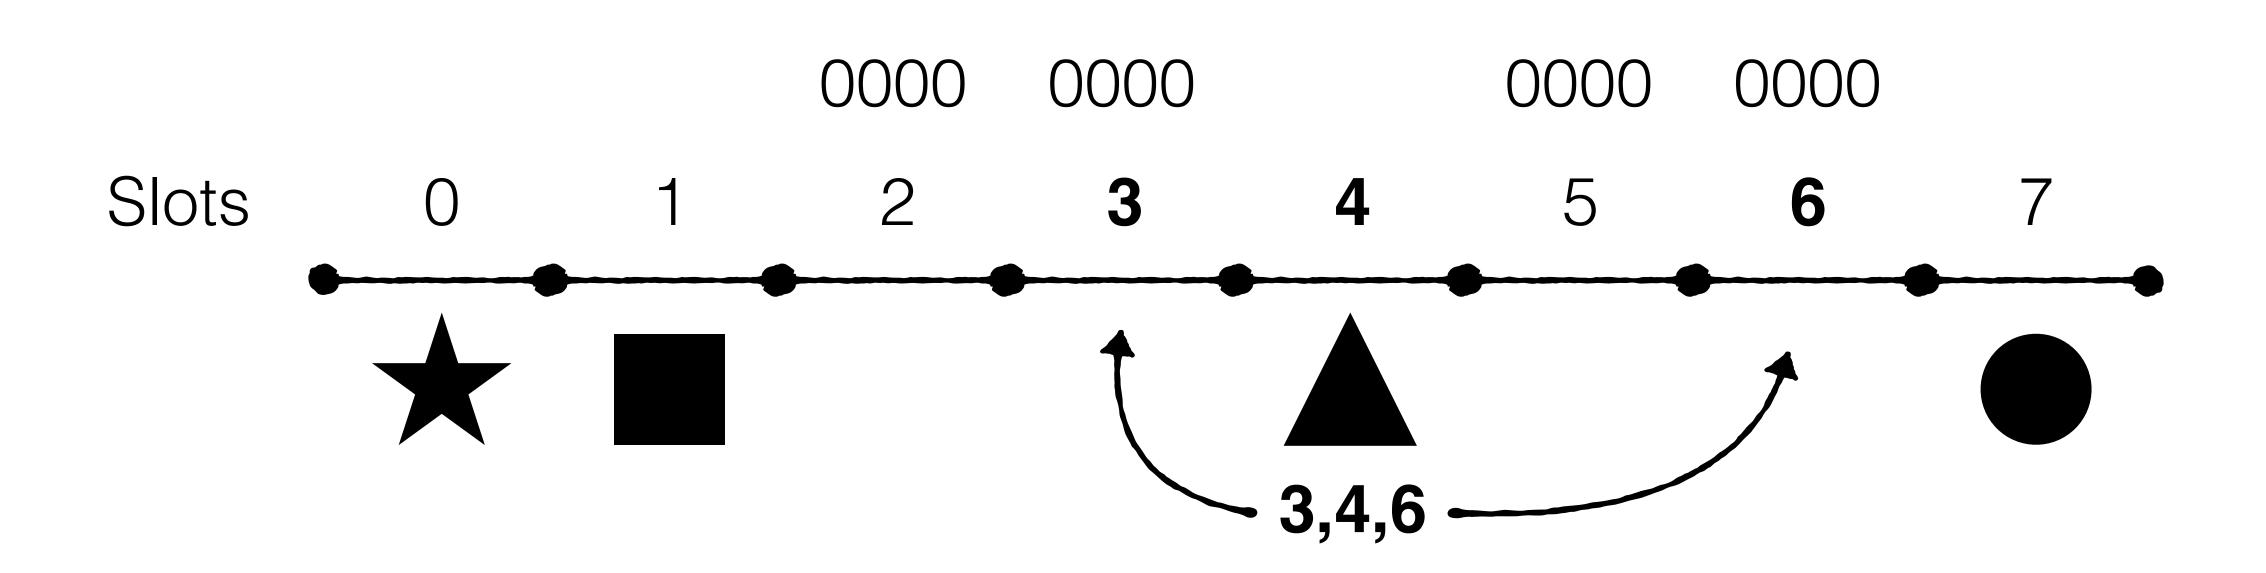
While not all keys have been assigned to a slot



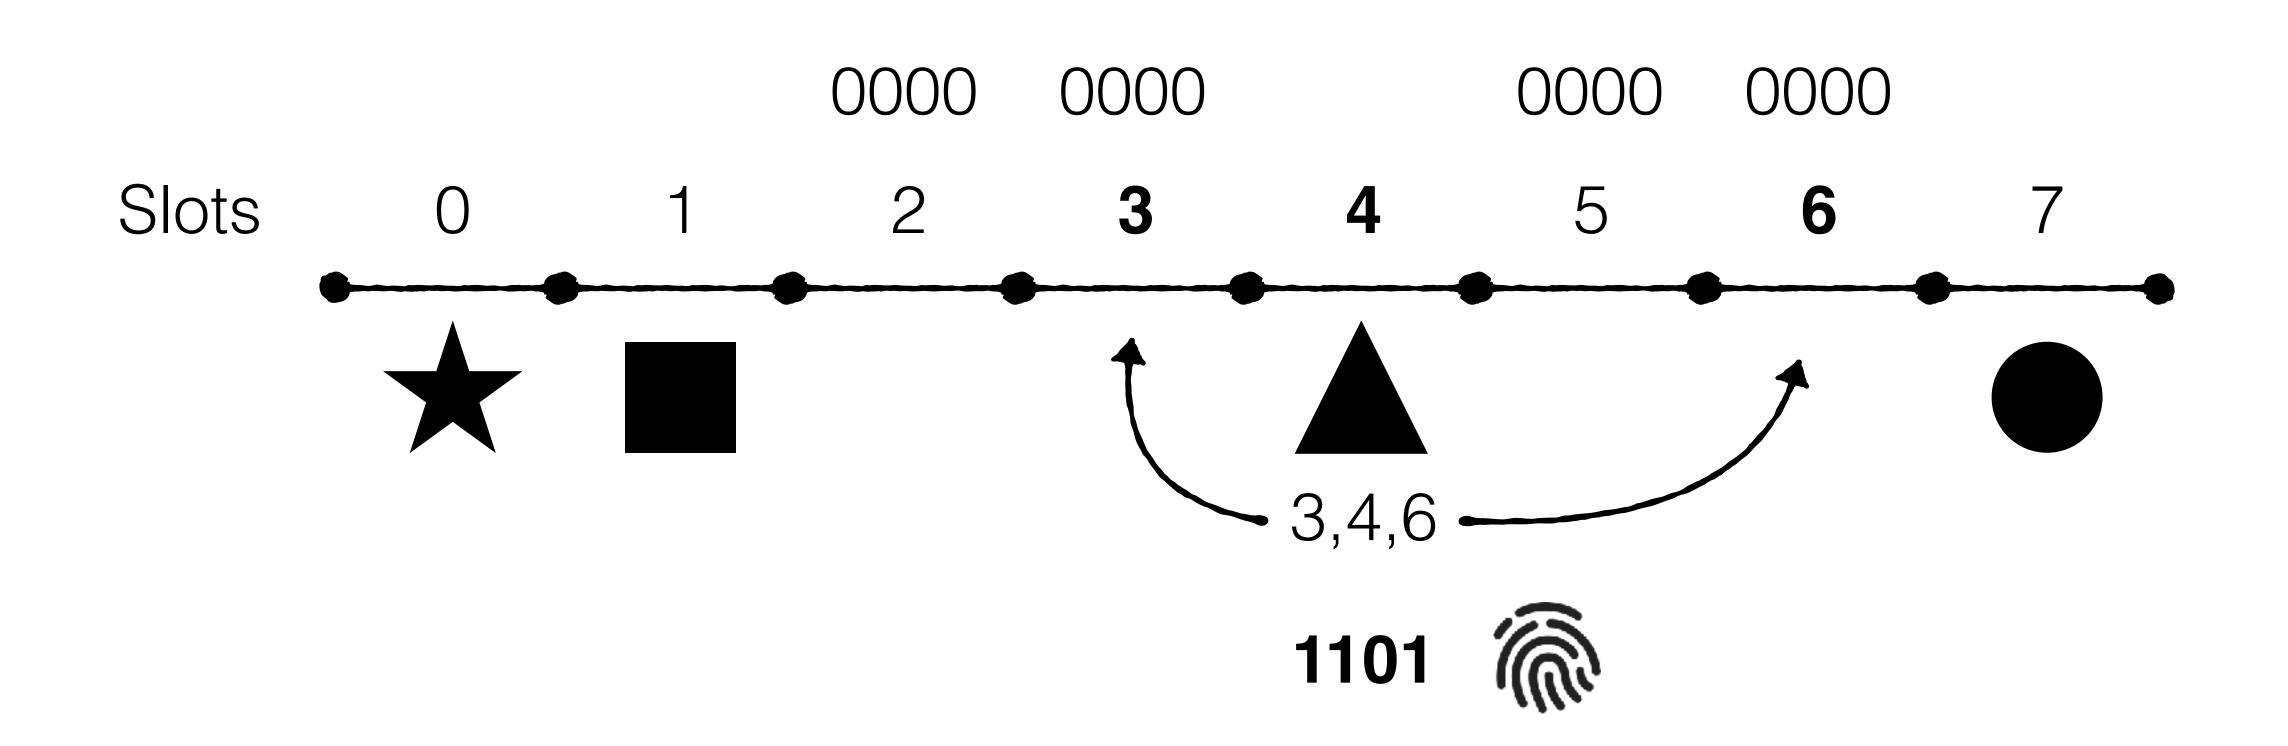




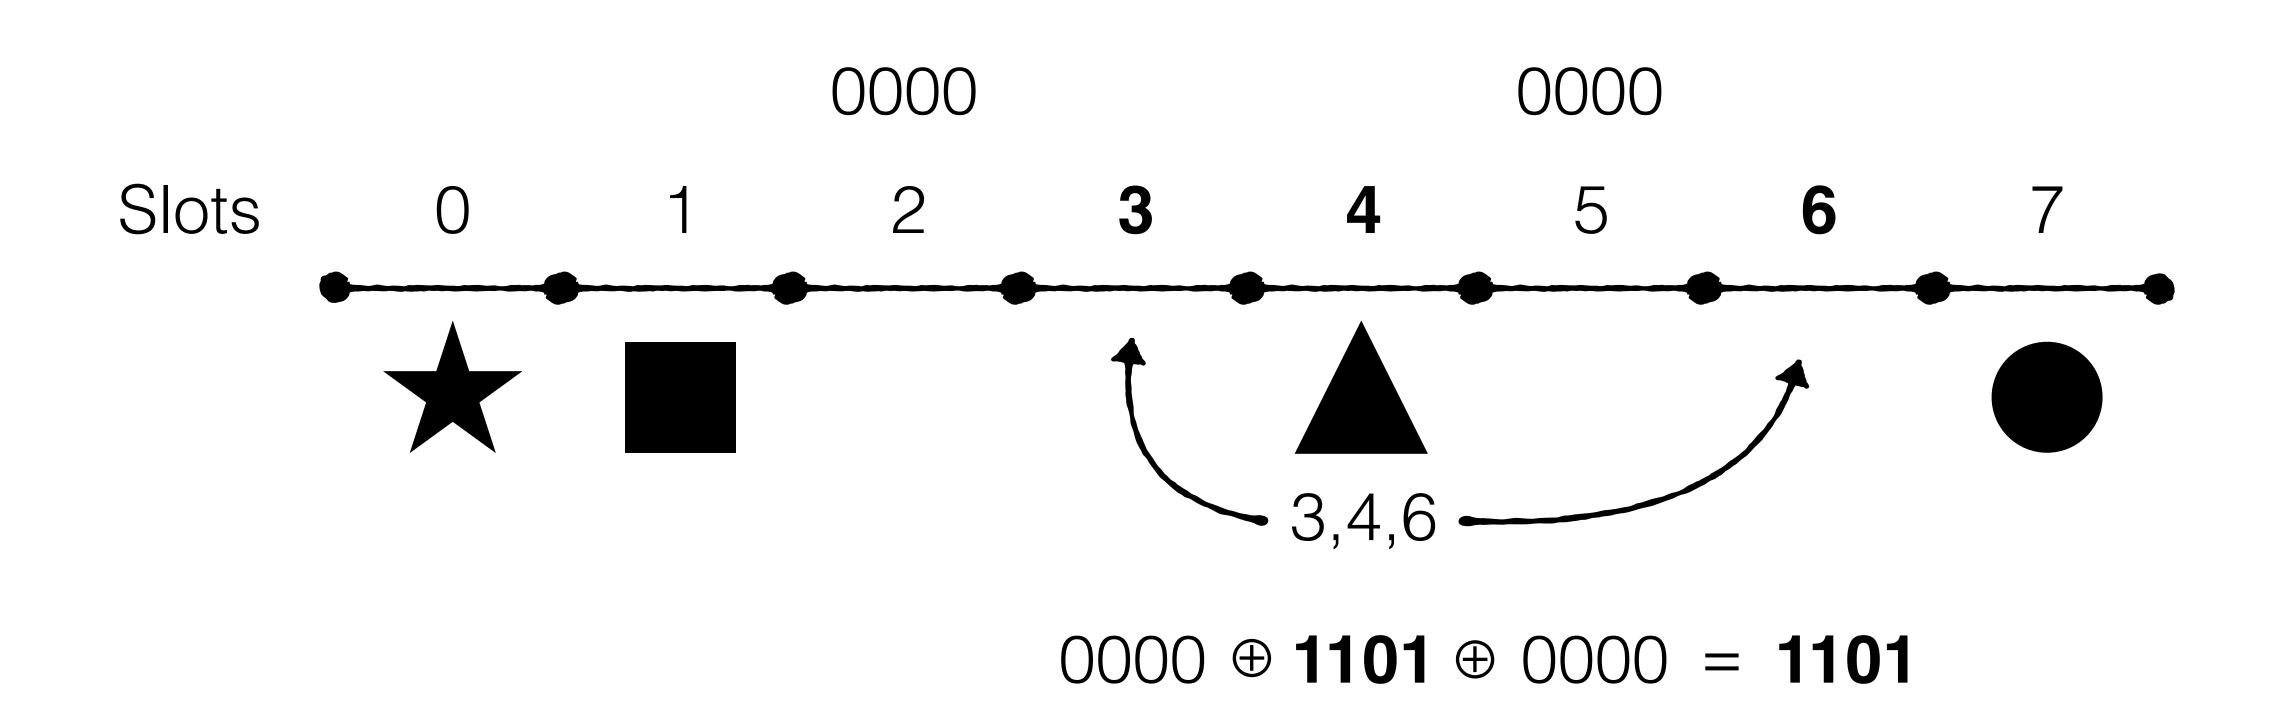
# Find some entry whose only other candidate slots are populated



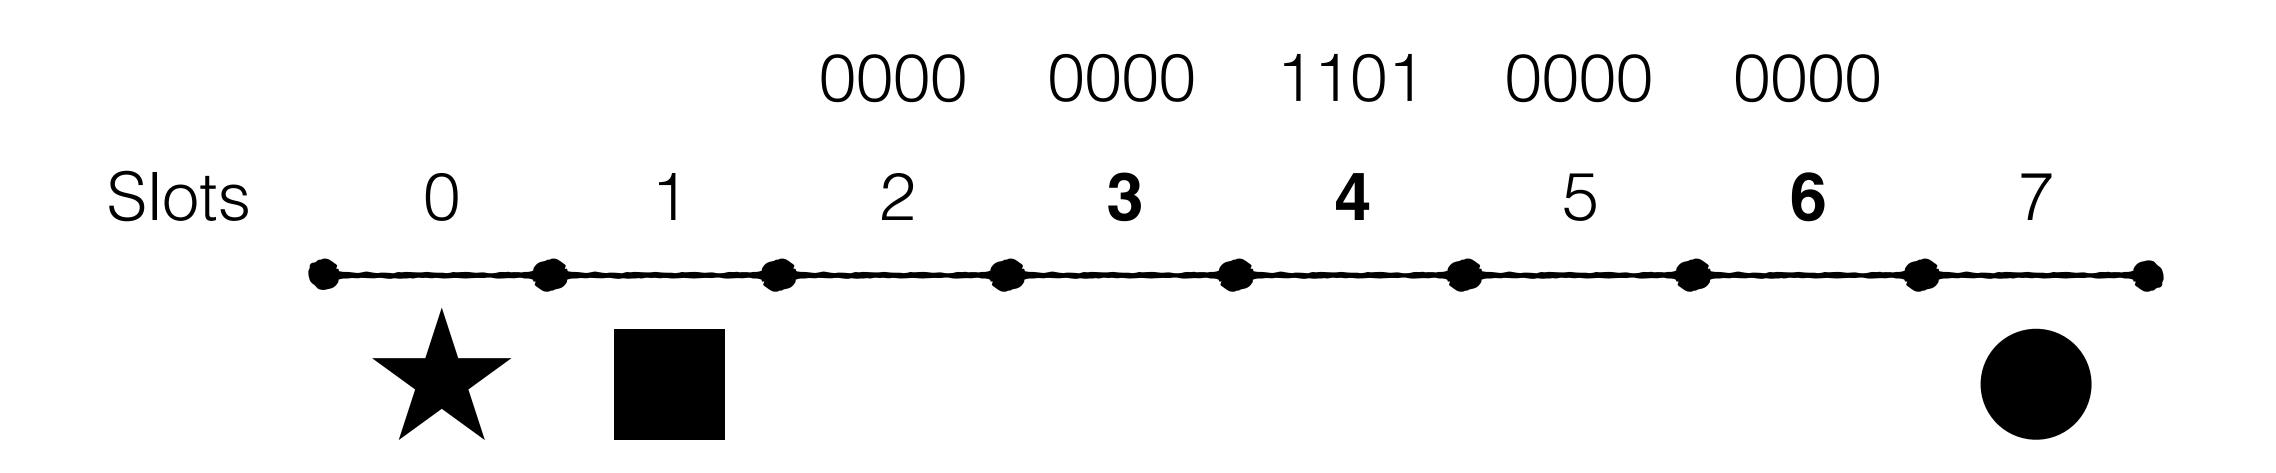
Find some entry whose only other candidate slots are filled



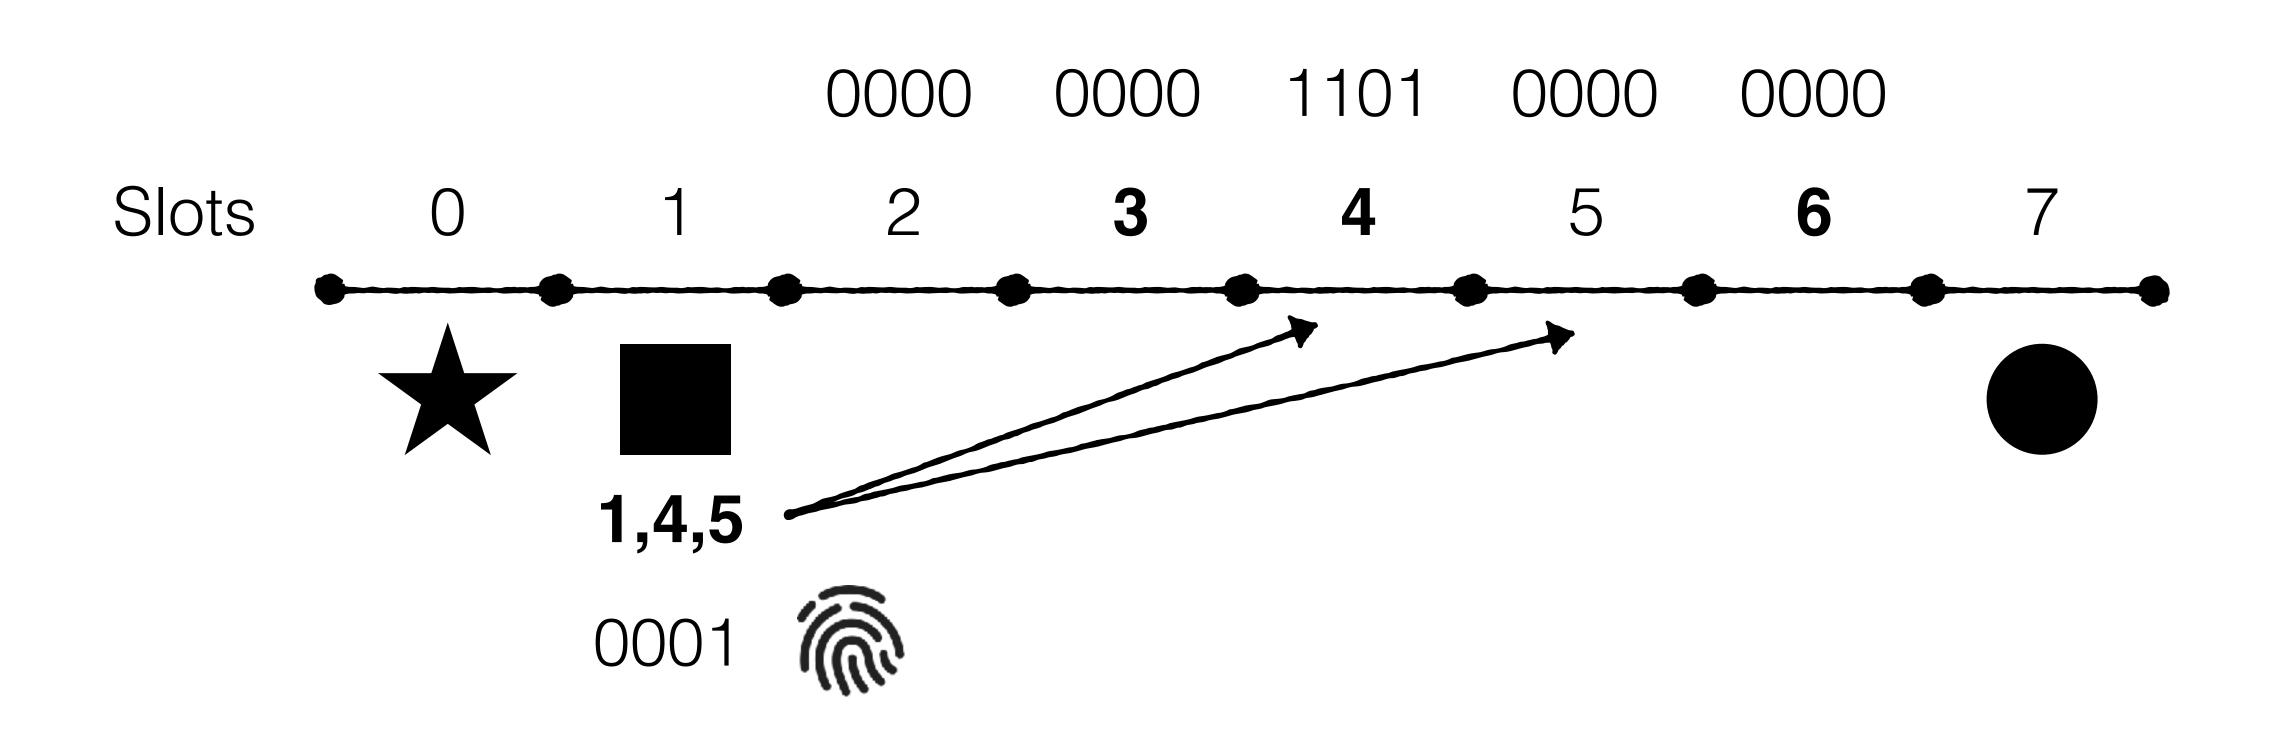
Find some entry whose only other candidate slots are filled



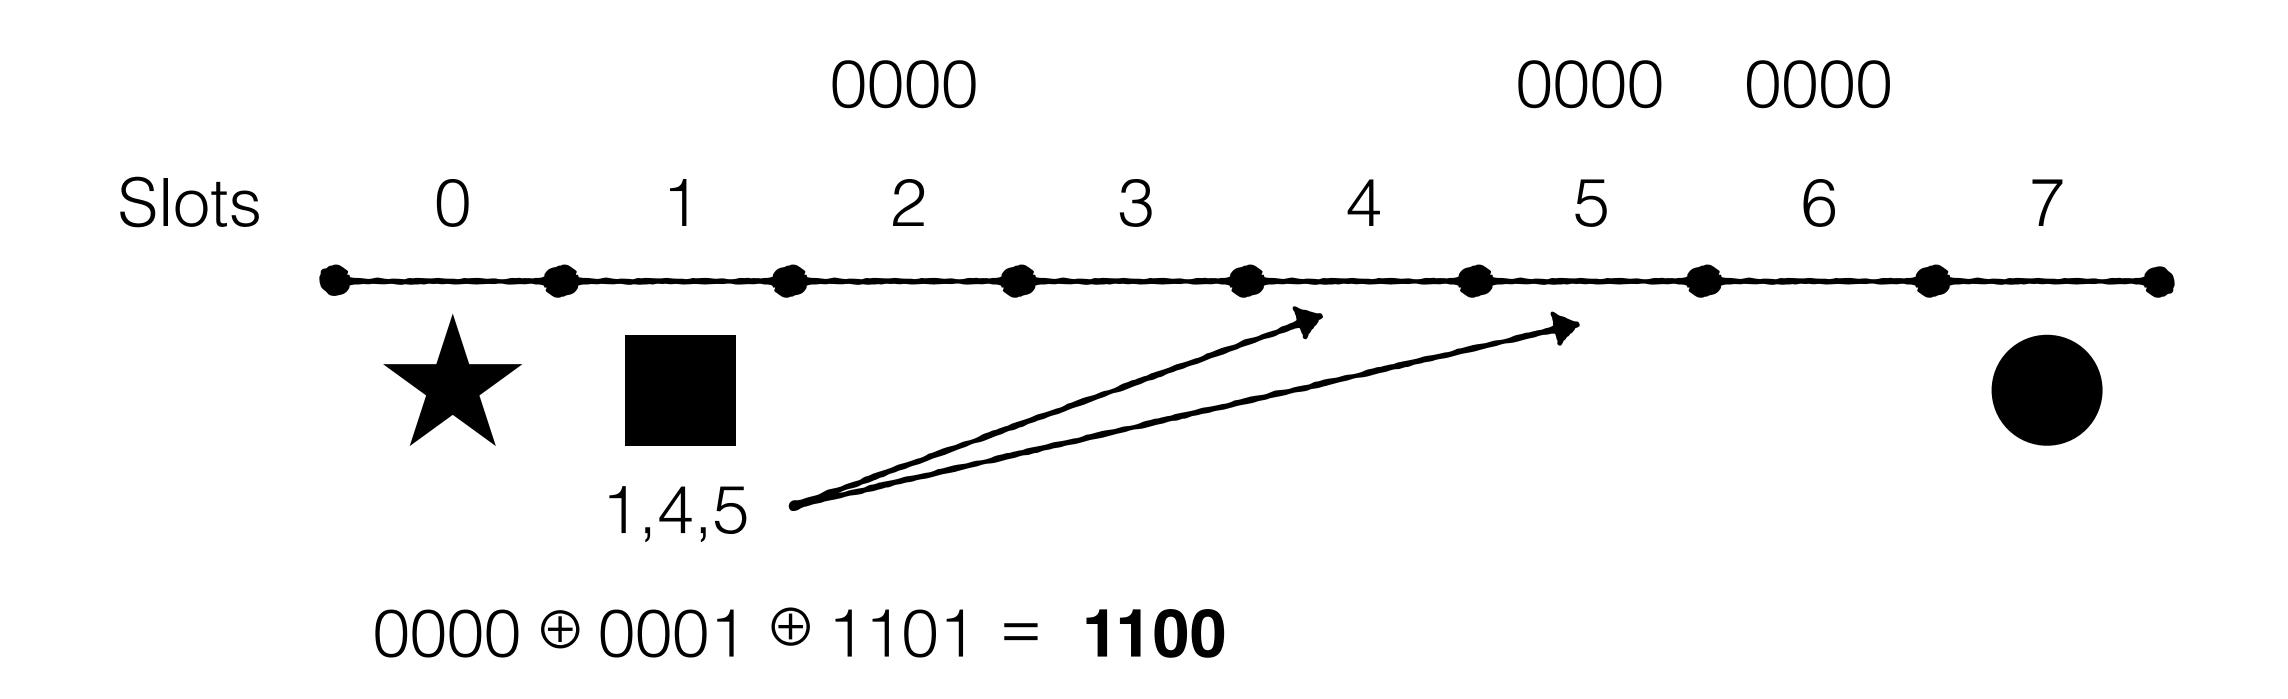
Find some entry whose only other candidate slots are filled



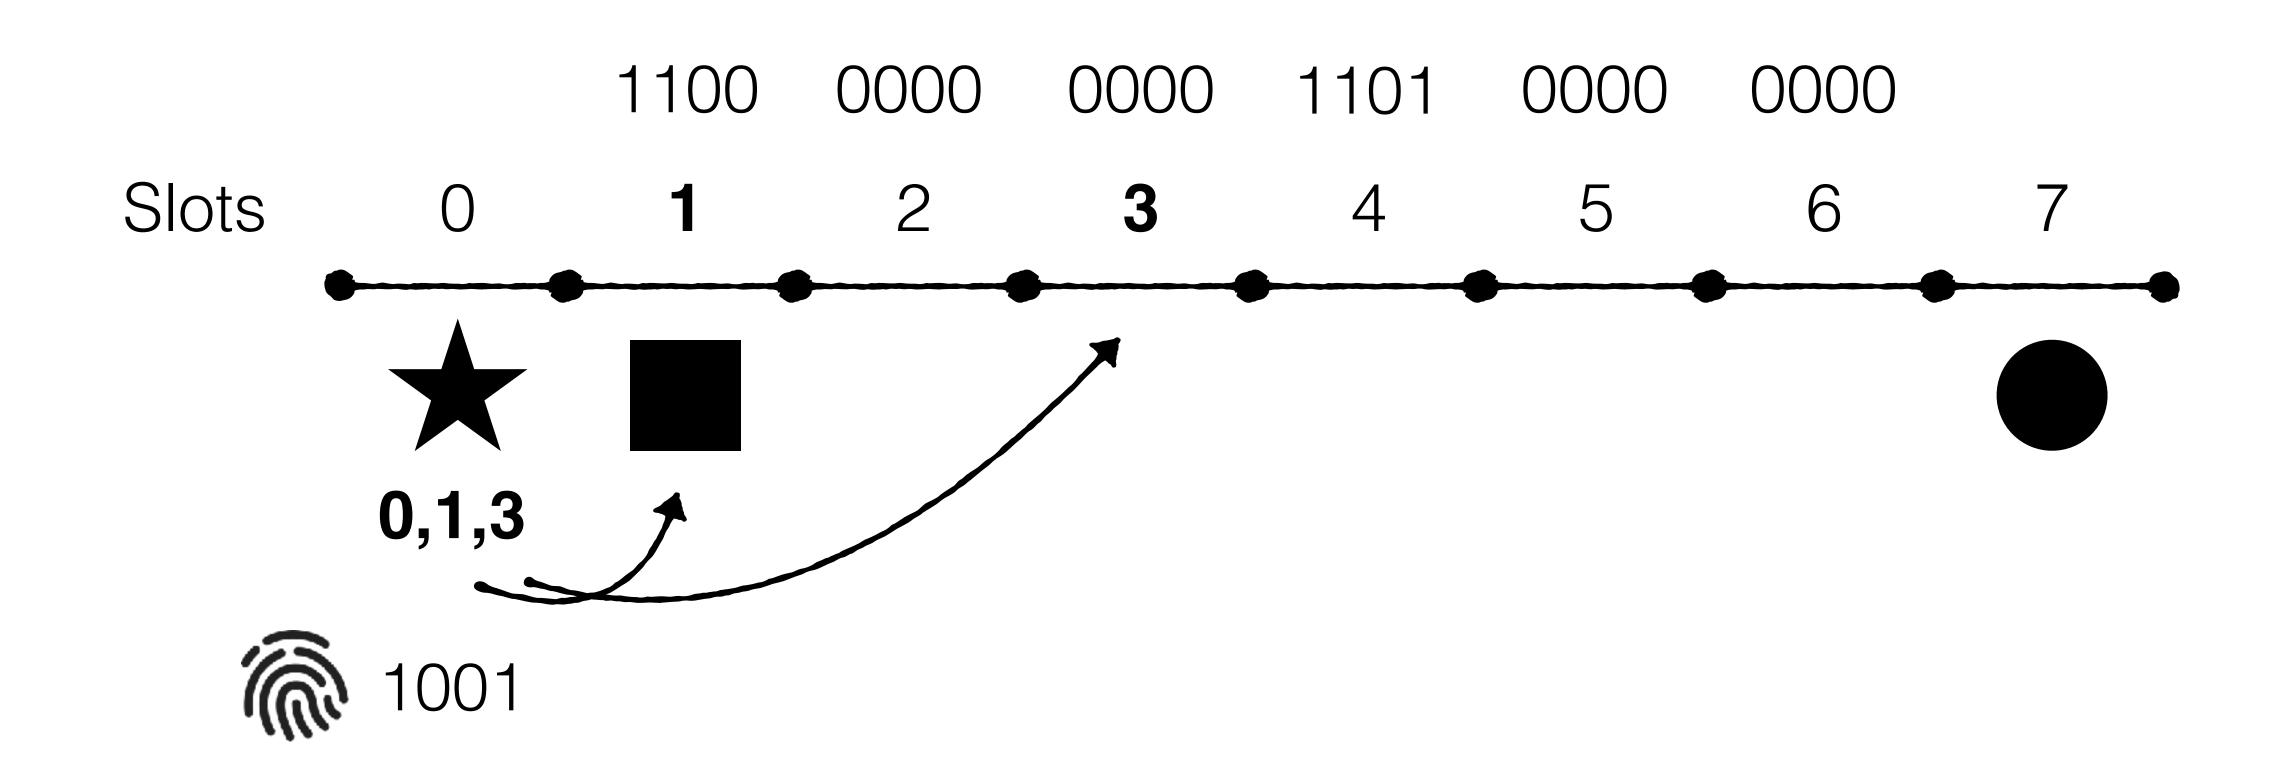
Find some entry whose only other candidate slots are filled



Find some entry whose only other candidate slots are filled

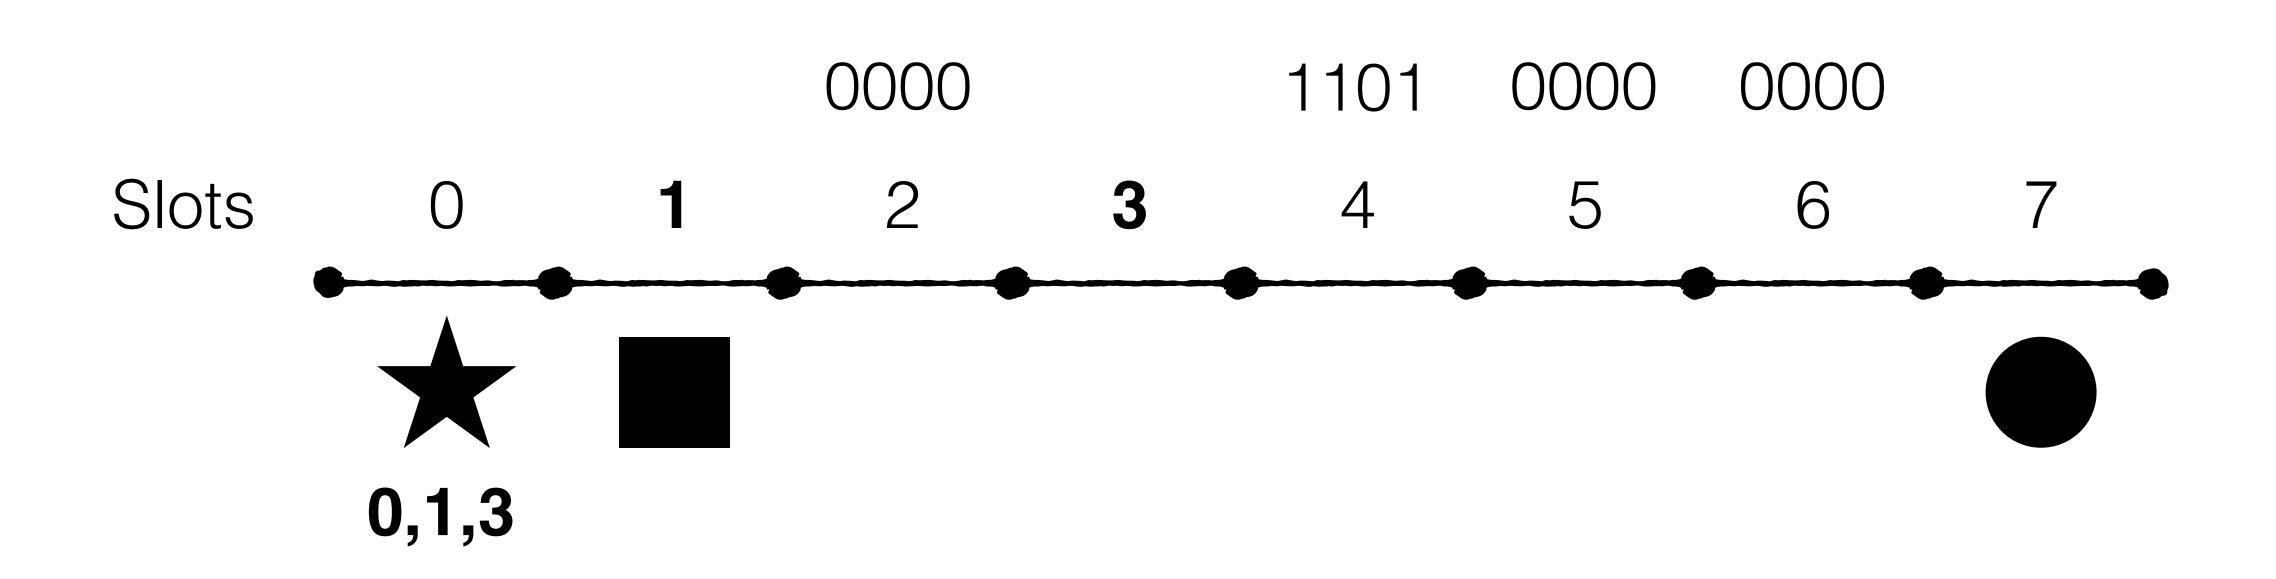


Find some entry whose only other candidate slots are filled



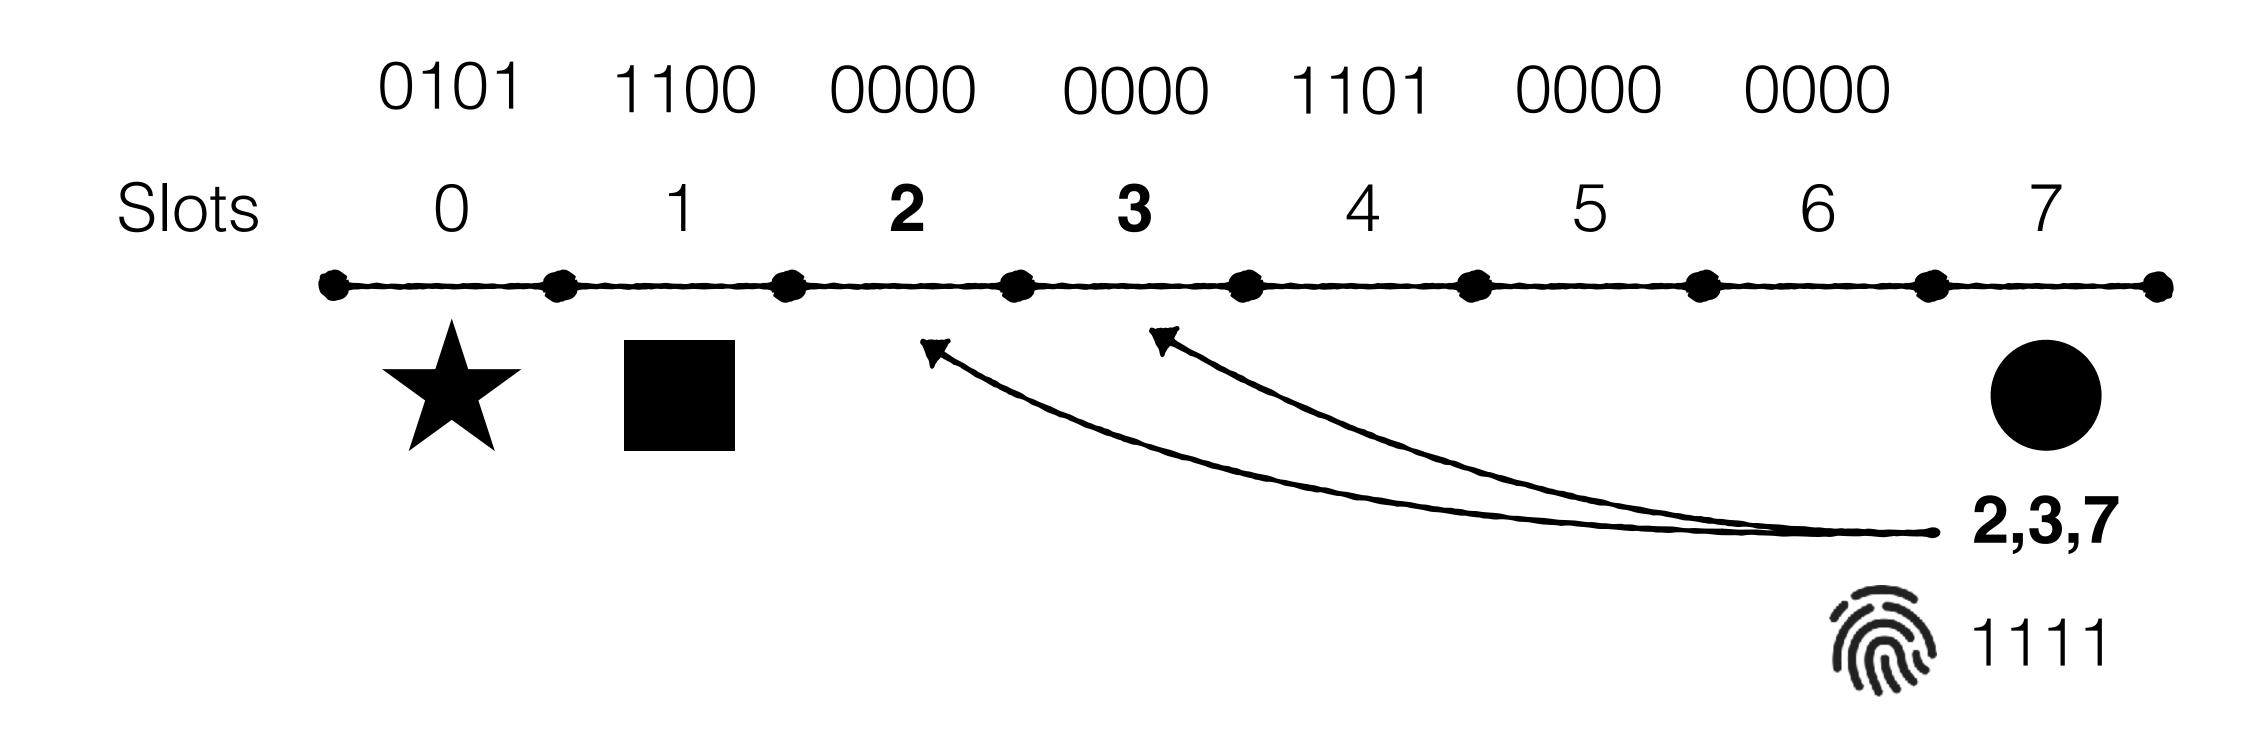
Find some entry whose only other candidate slots are filled

Xor its fingerprint with content at other slots and store

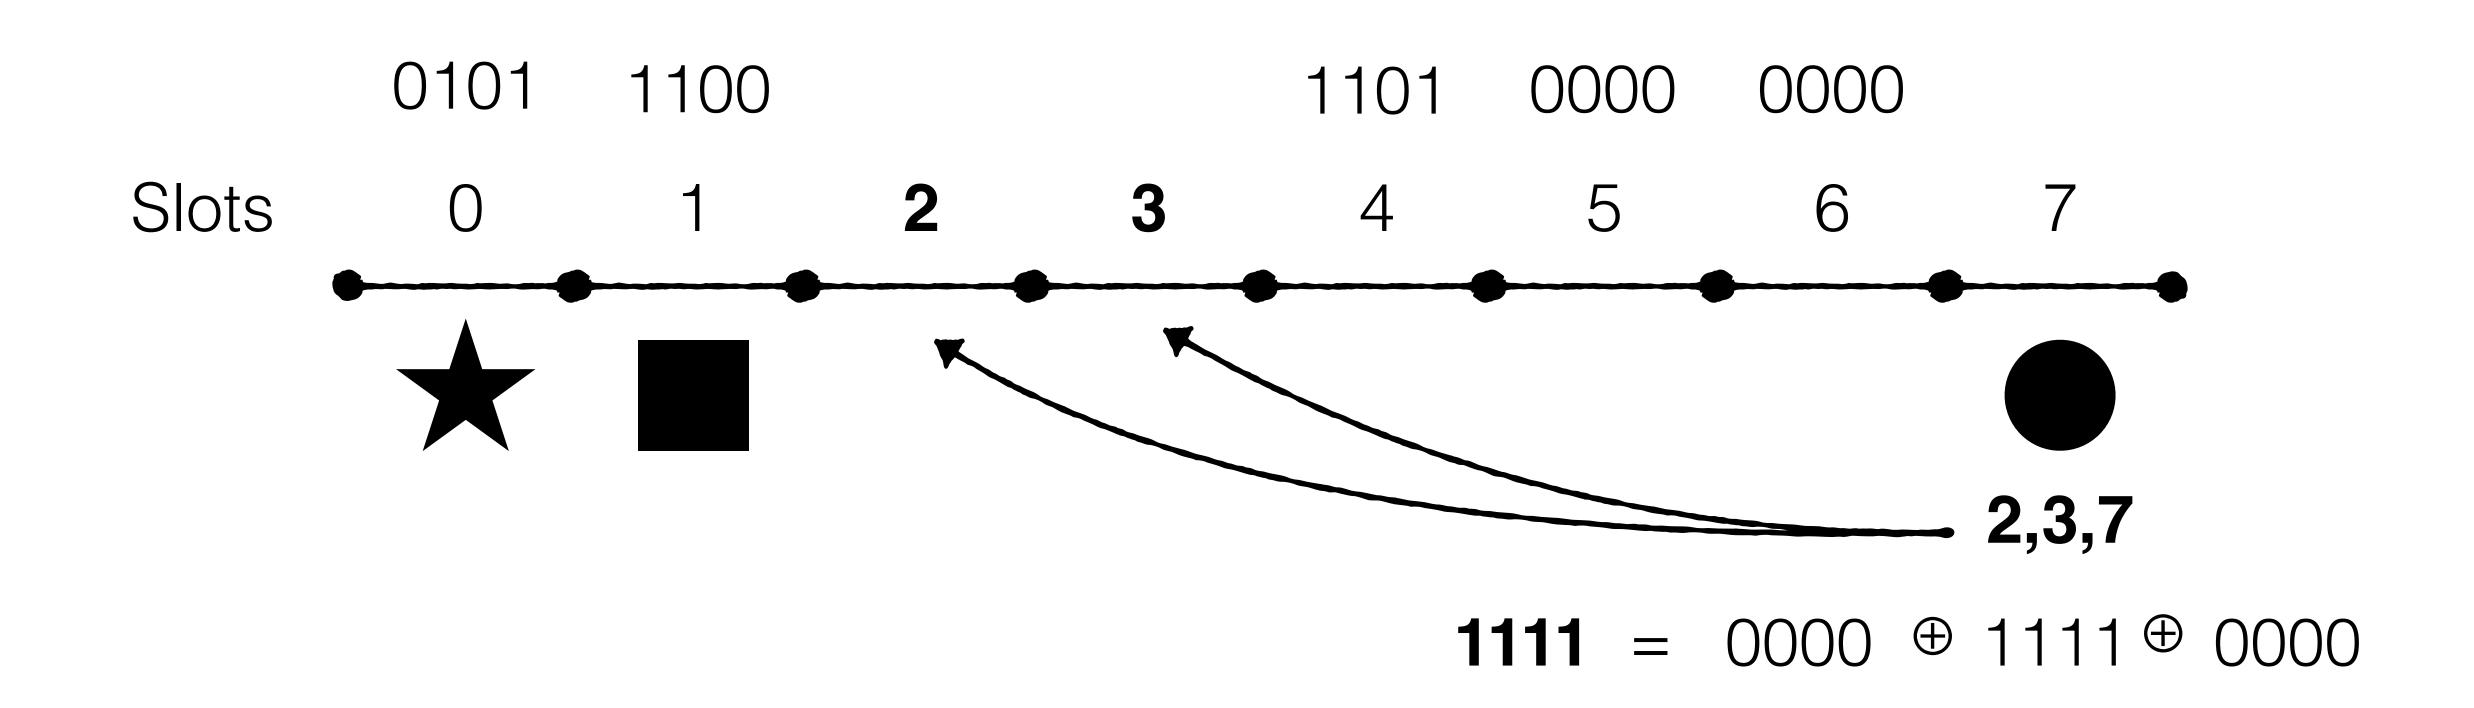


1100 ⊕ 1001 ⊕ 0000 = **0101** 

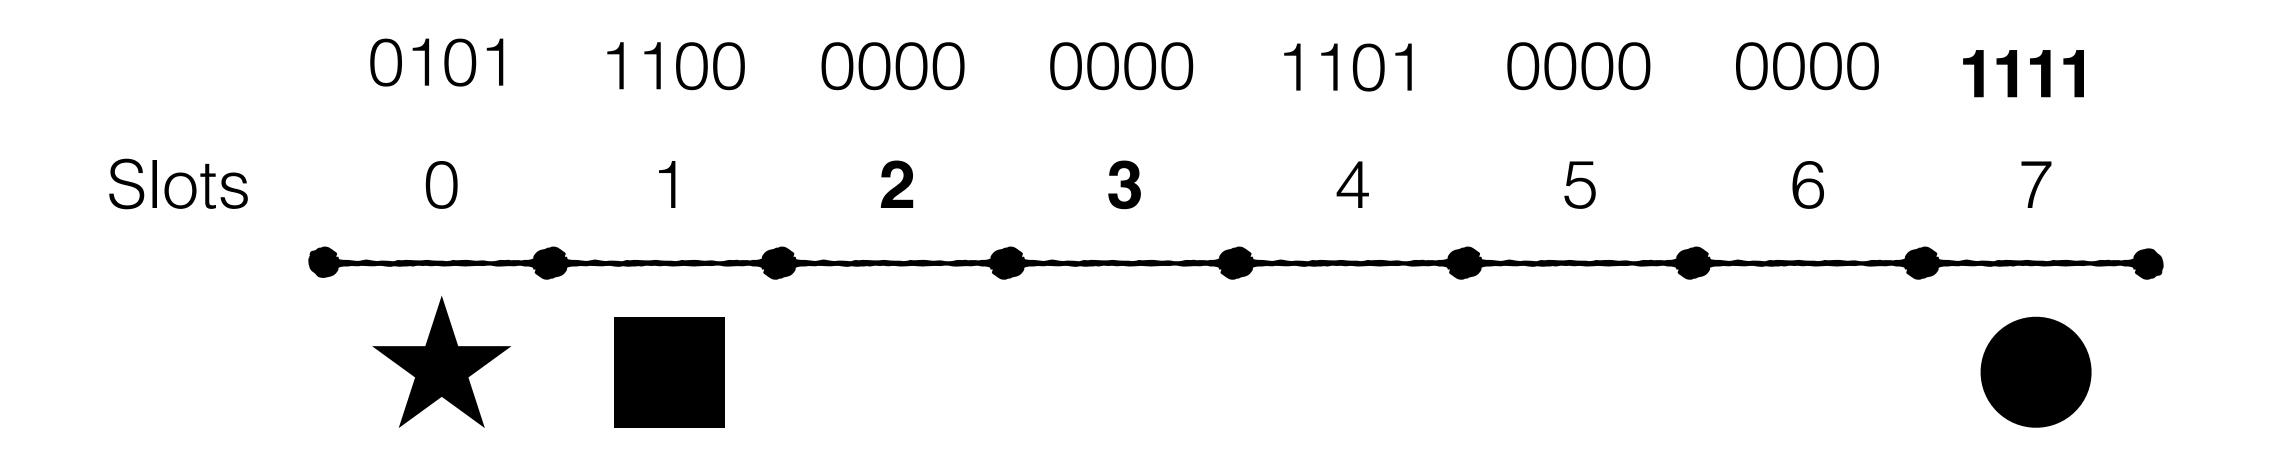
Find some entry whose only other candidate slots are filled



Find some entry whose only other candidate slots are filled



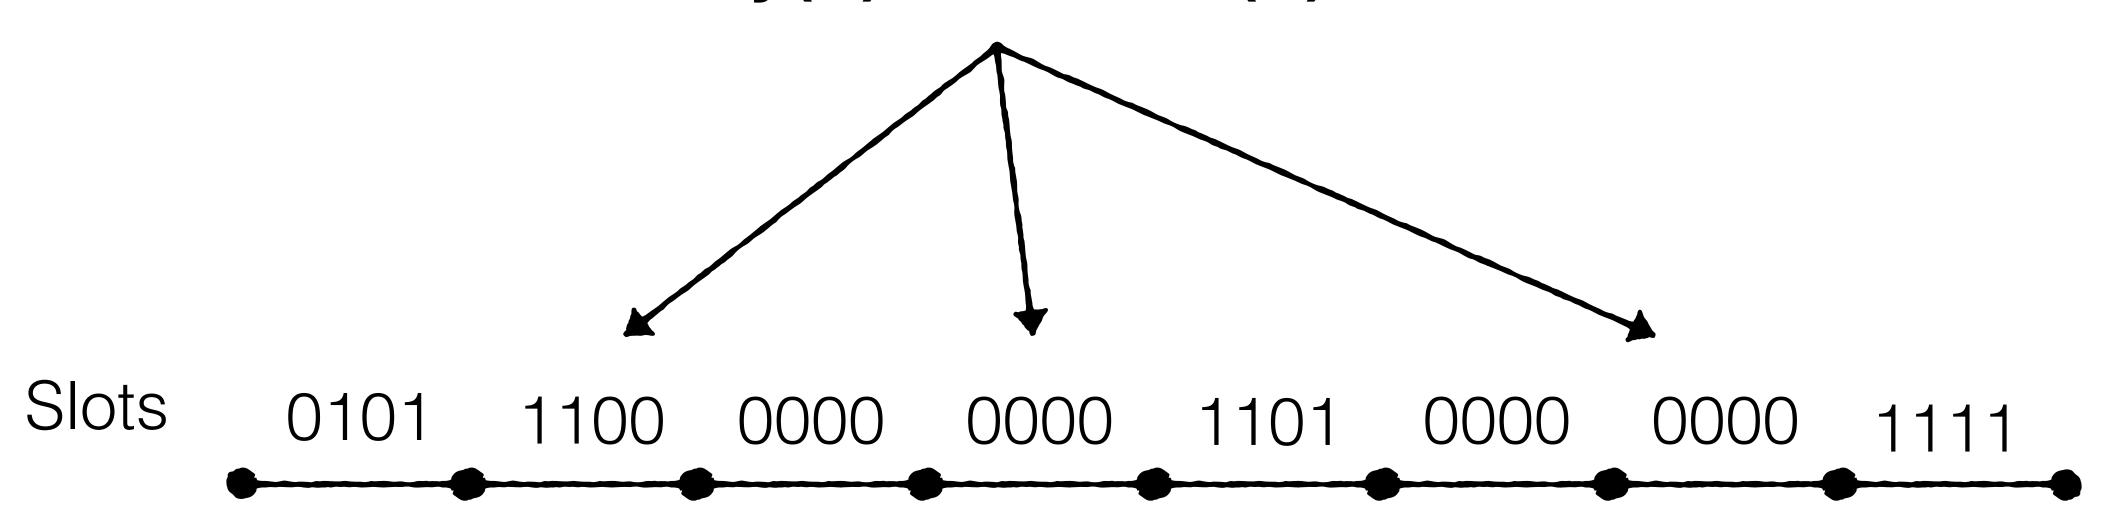
Find some entry whose only other candidate slots are filled



We're done:)



# Query(X) where FP(X) = 1001



### Query(X) where FP(X) = 1001

 $1100 \oplus 0000 \oplus 0000 = 1100$ 



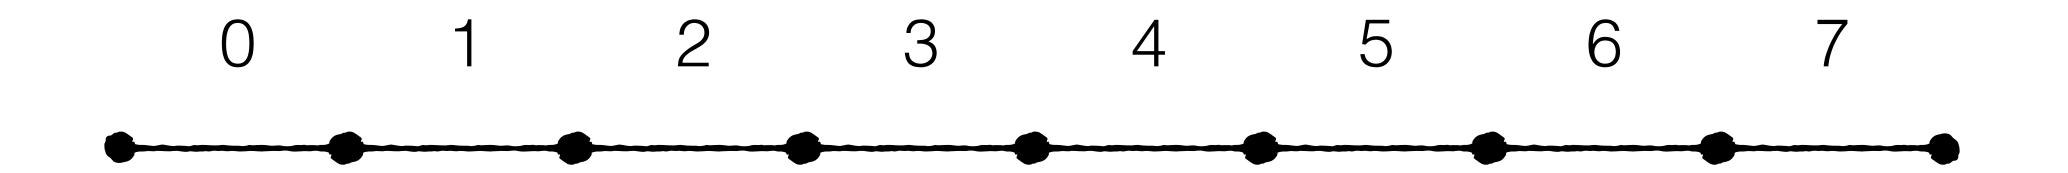
#### Not a fingerprint match so return negative

Query(X) where FP(X) = 1001

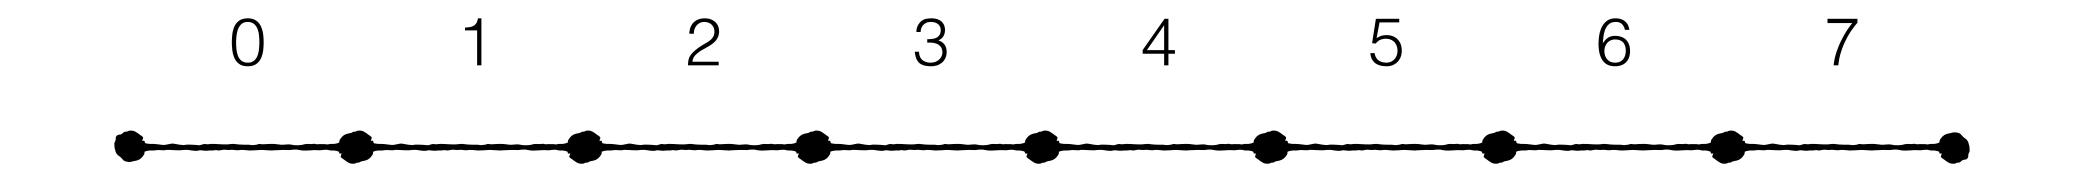
 $1100 \oplus 0000 \oplus 0000 = 1100$ 



# Construction can fail if there is no entry we can peel

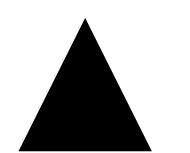


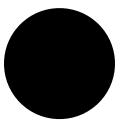
# Construction can fail if there is no entry we can peel



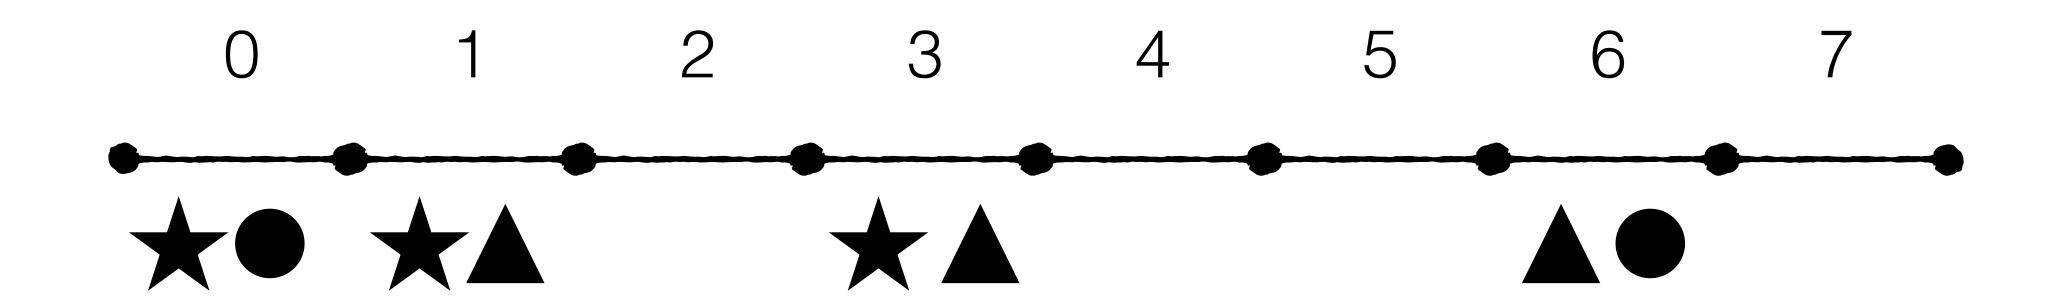
**Example:** 





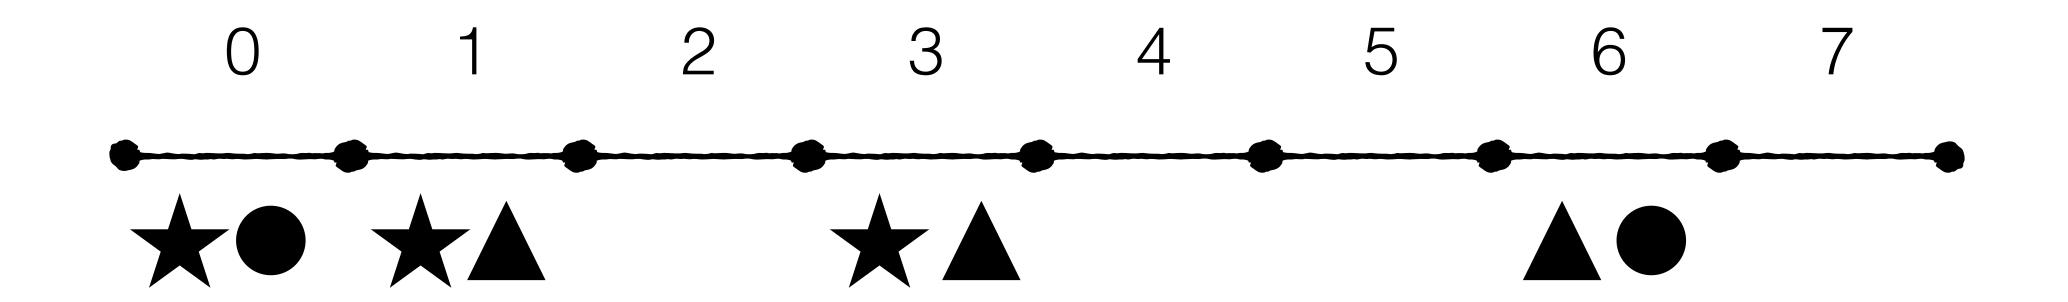


Construction can fail if there is no entry we can peel



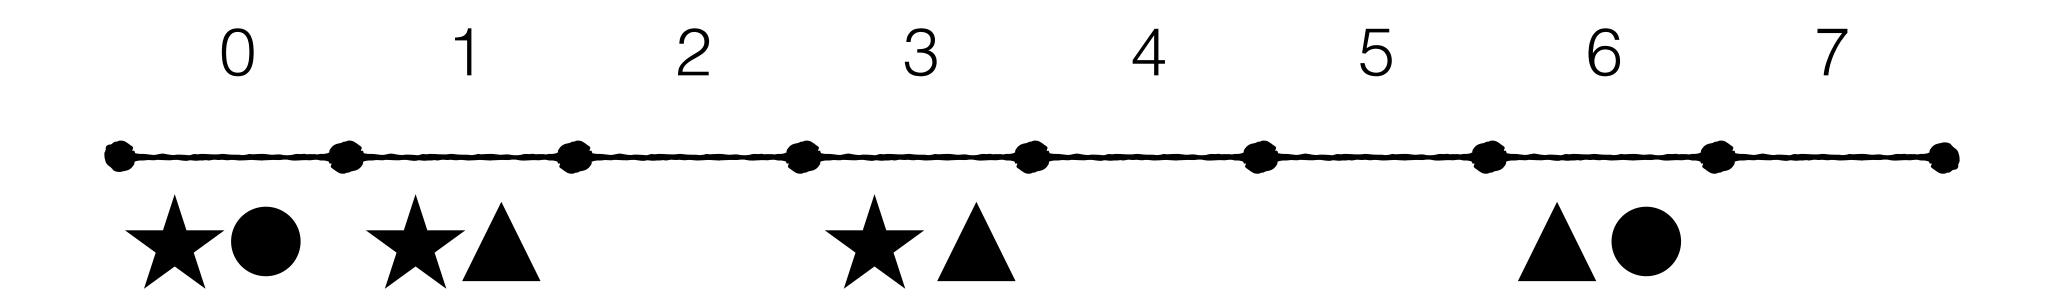
No slot has one entry uniquely mapping to it

If we fail, we must restart from scratch.



If we fail, we must restart from scratch.

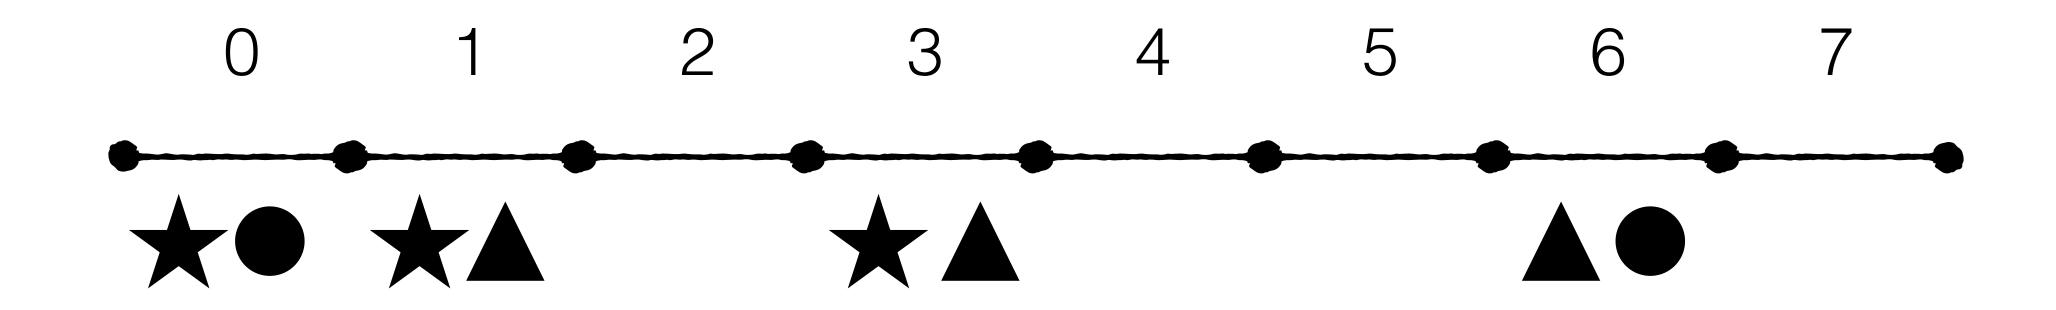
# free space is necessary to succeed with high probability



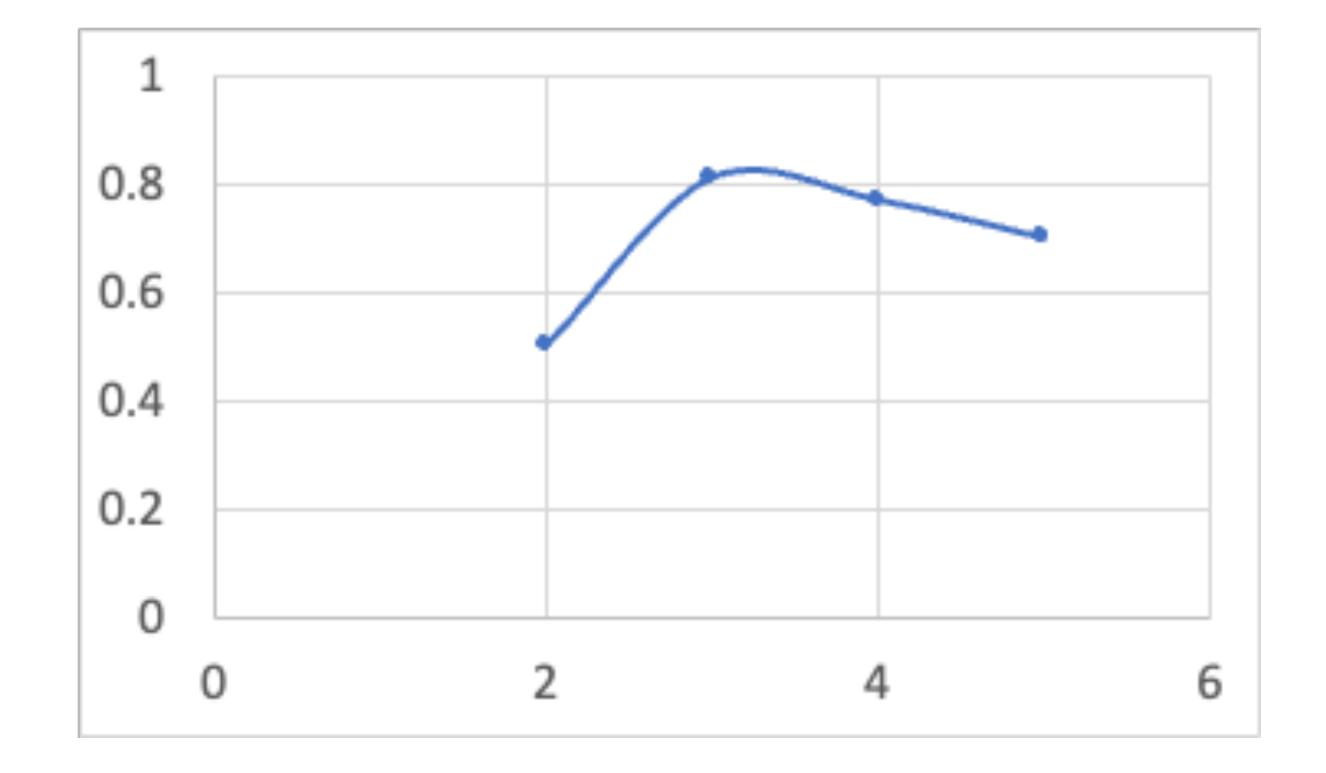
If we fail, we must restart from scratch.

free space is necessary to succeed with high probability

# What's the interplay between free space and # number of hash functions?



# Utilization

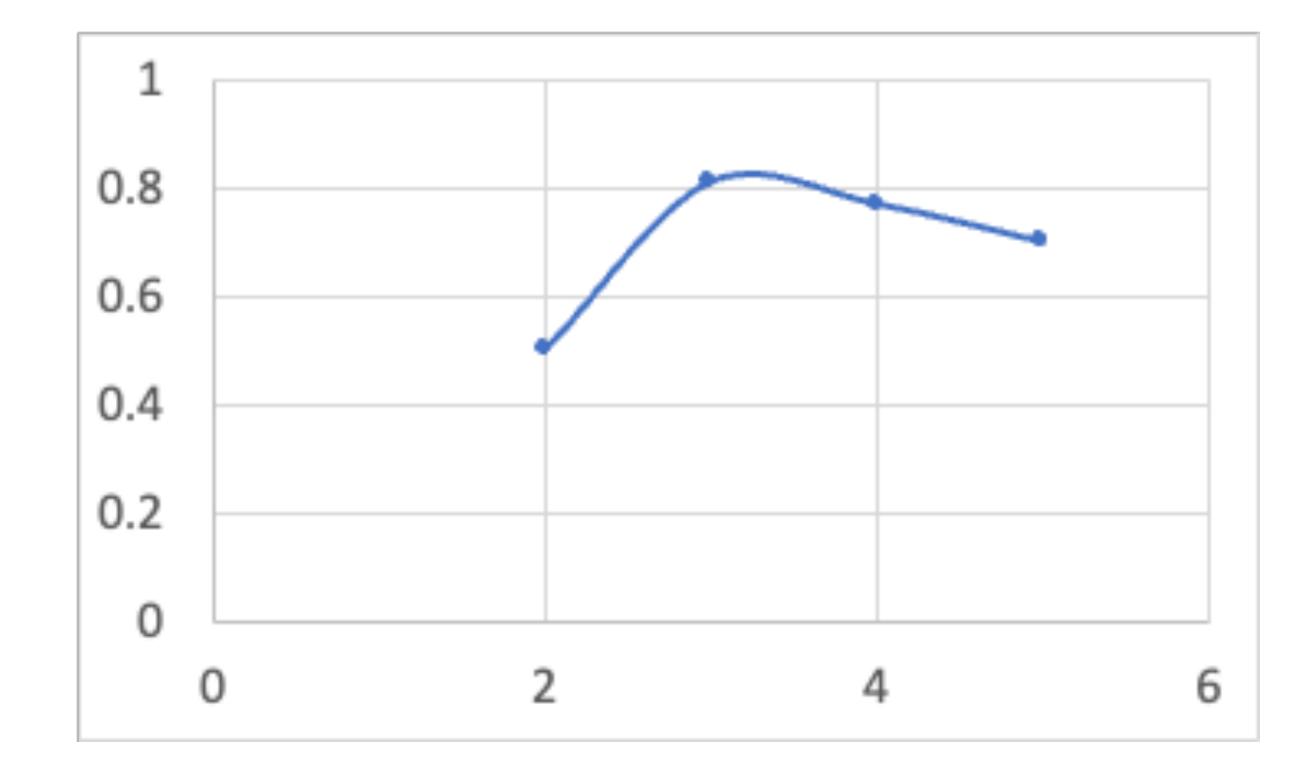


# hash functions

# Not enough placement flexibility





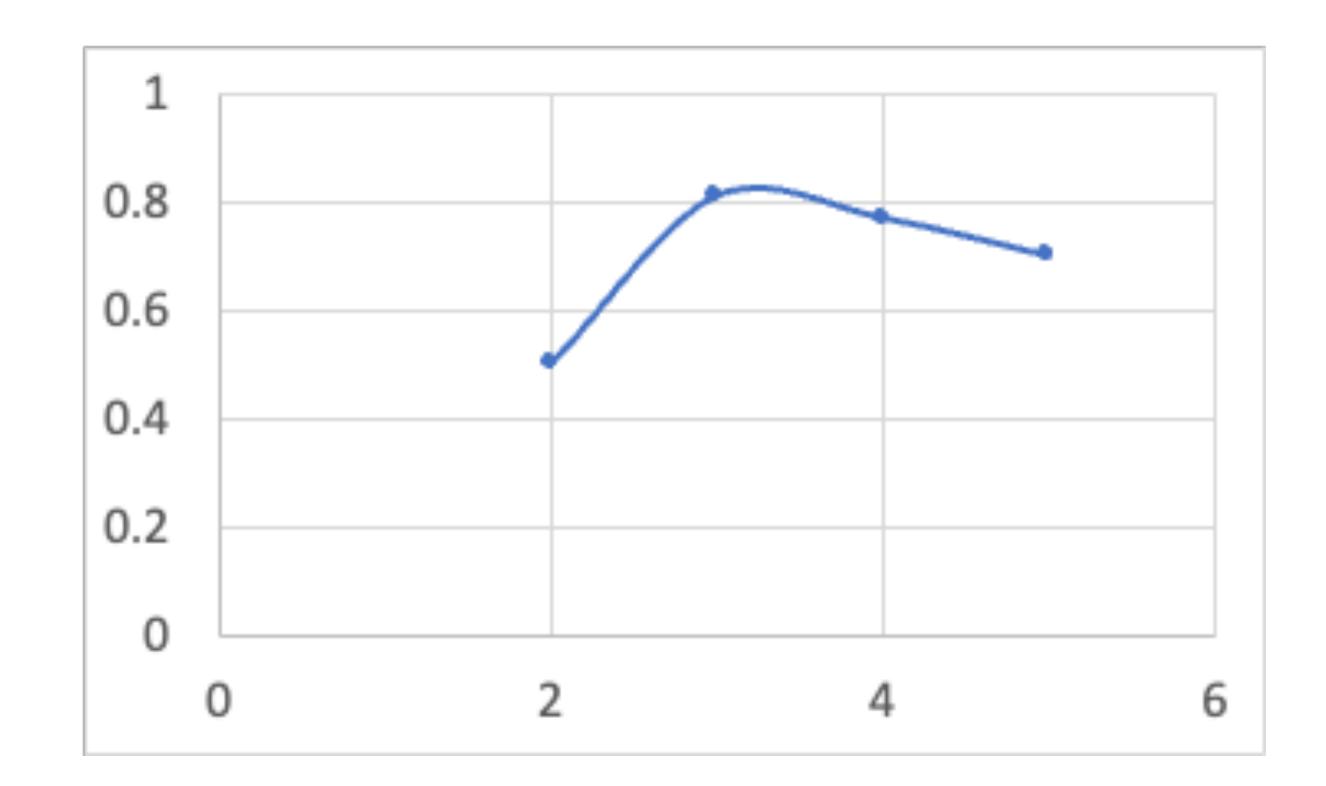


# hash functions

# Too many items hashing to each slot





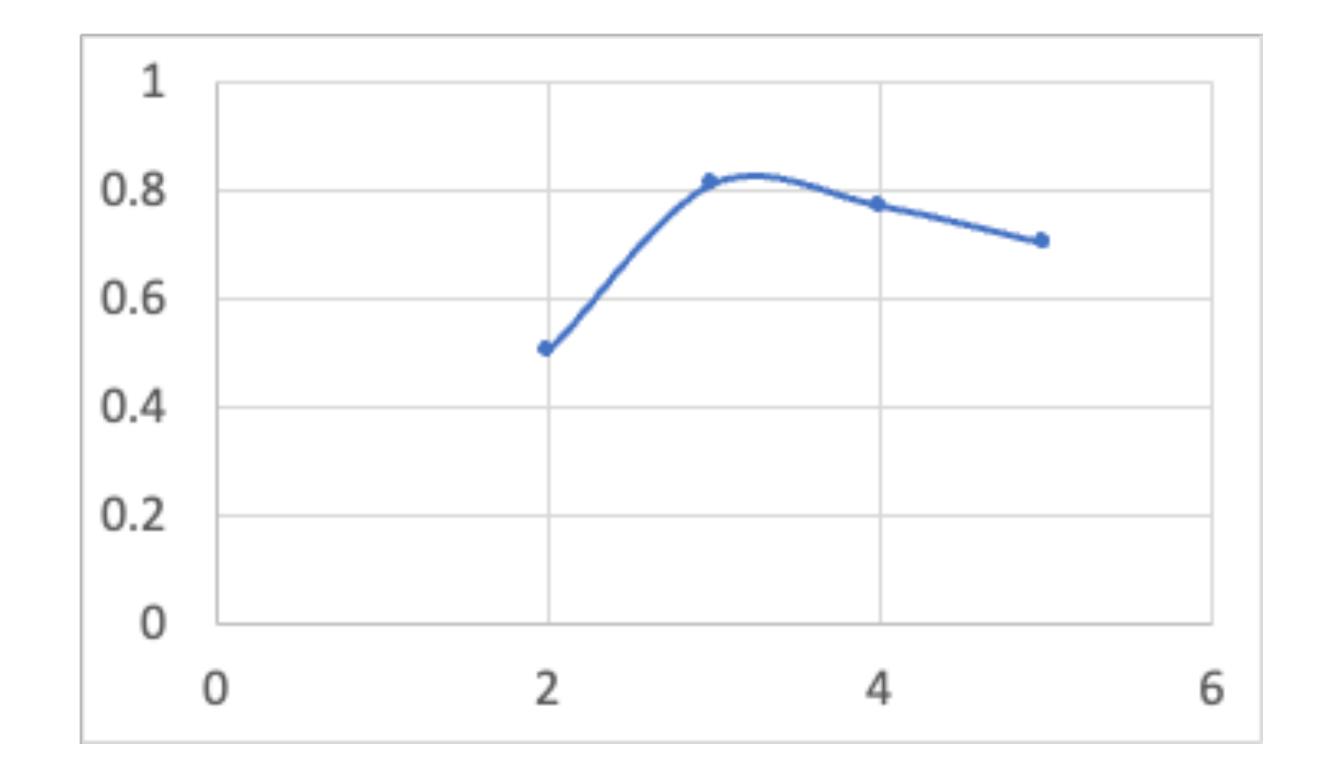


# hash functions

# Optimal 0.81



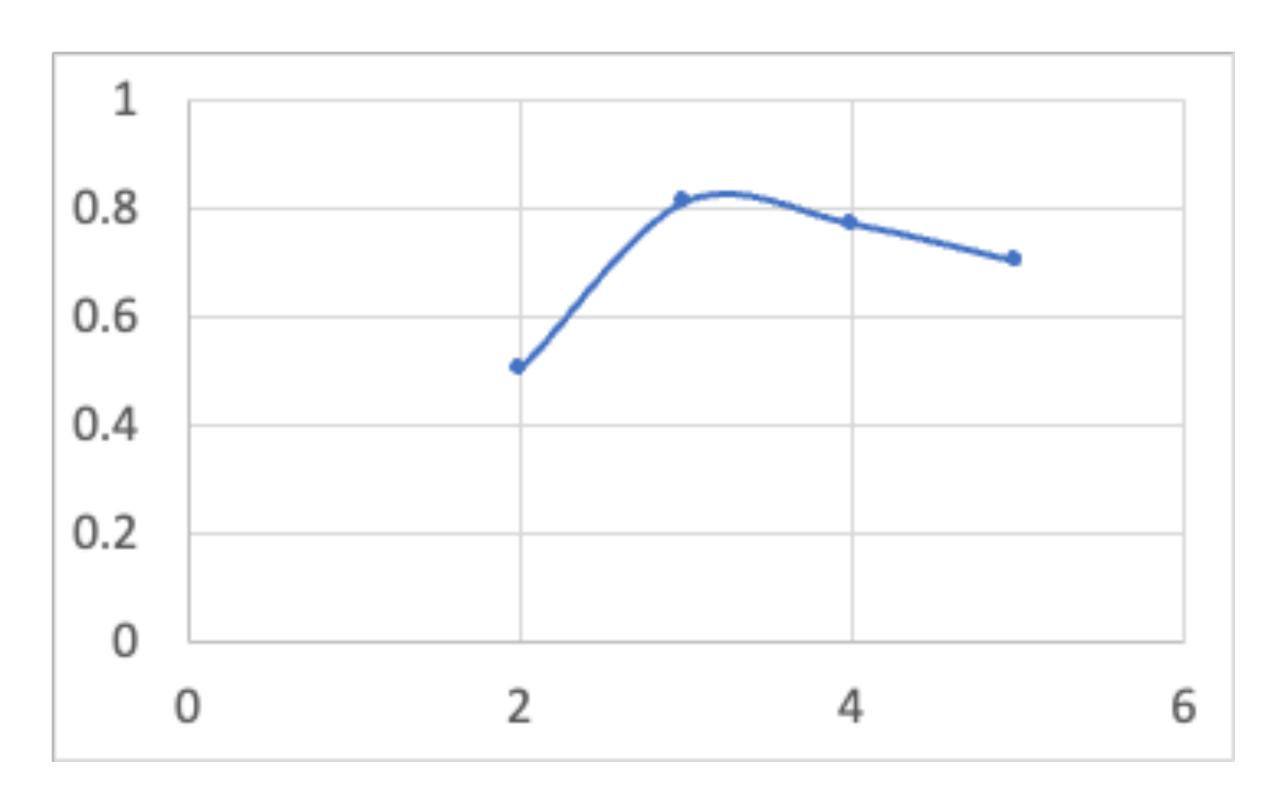




# hash functions

# Optimal 0.81





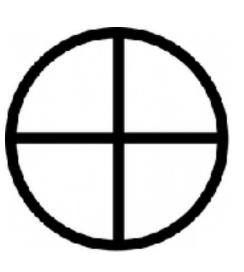
Similar to finding the optimal # hash functions with Bloom filters:)

XOR

Idealized



 $\approx 2 - M/N \cdot 0.69$ 





≈ 2 -M/N · 0.81

 $\approx 2 - M/N$ 

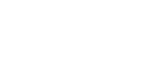


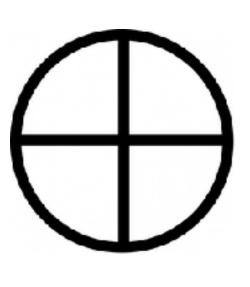


Idealized



 $\approx 2 - M/N \cdot 0.69$ 





 $\approx 2 - M/N \cdot 0.81$ 



≈ 2 -M/N · 0.92



 $\approx 2 - M/N$ 

Denser XOR filter

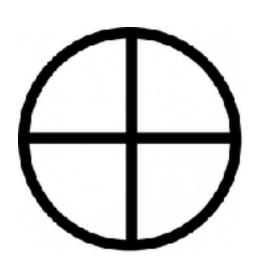
XOR

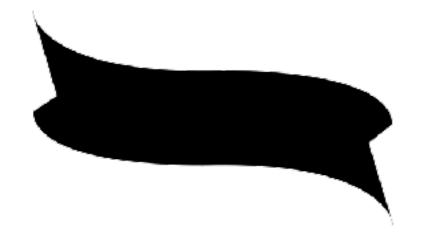
Ribbon

Idealized



 $\approx 2 - M/N \cdot 0.69$ 







 $\approx 2 - M/N \cdot 0.81$ 

≈ 2 -M/N · 0.92

 $\approx 2 - M/N$ 

Denser XOR filter

In RocksDB since 2020

XOR

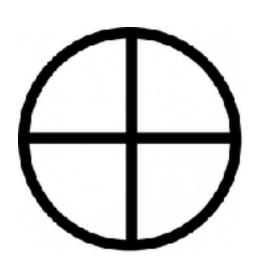
Ribbon

XOR filter w.

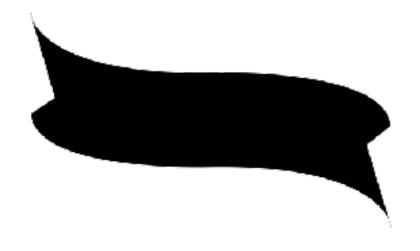
Spatial Coupling



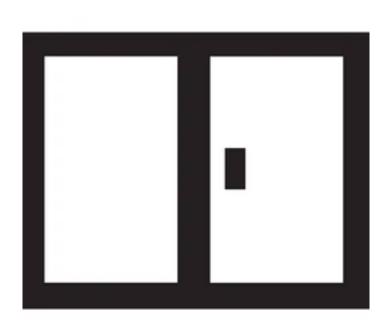




 $\approx 2 - M/N \cdot 0.81$ 



 $\approx 2 - M/N \cdot 0.92$ 



 $\approx 2 - M/N$ 

Approach ideal



Construction =



Positive Query =



Avg. Negative Query =



Construction = O(N)



Positive Query =



Avg. Negative Query =



Construction = O(N)



Positive Query = 3



Avg. Negative Query = 3



Construction = O(N)



Positive Query = 3



Avg. Negative Query = 3

Not as good as blocked Bloom filters

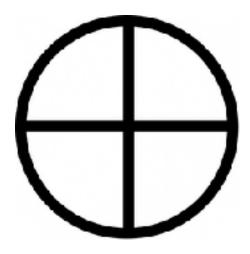
Blocked Bloom

XOR

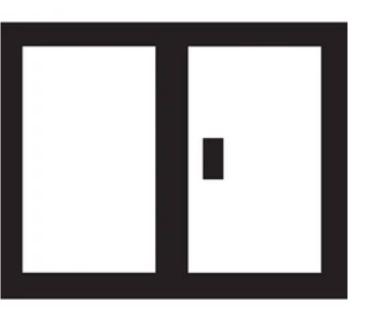
Ribbon

Spatial Coupling









Faster

Lower FPR

And now: office hours:)