


# The Memory Hierarchy & Bloom Filters

**Niv Dayan**



# Memory Hierarchy

CPU caches



memory



SSD




disk




# Memory Hierarchy

CPU caches



memory



SSD



disk




**Fast & Expensive**



**Slow & Cheap**


## Latency

CPU caches



**10 ns**

memory



**100 ns**

SSD




**100  $\mu$ s**

disk



**10 ms**

# Latency



10 ns



100 ns





100  $\mu$ s





10 ms






**Most data is here**



**Sometimes, we must fetch  
data from storage**



get key X



X

Sometimes, we must fetch  
data from storage





**But sometimes, user search for non-existing data**



**get key Q**



**{X, Y, Z}**


But sometimes, user search for non-existing data



e.g., does key Q exist?




{X, Y, Z}





**Can we quickly tell if a key does not exist so  
that we don't have to search storage?**

# What is a Filter?




# What is a Filter?

Does X exist?





Set




# What is a Filter?


Does X exist?



Set



**No false  
negatives**



**false positives with  
tunable probability**





**Does key X exist**



Data






**Does key  
X exist**



**Memory**



**Data**






# If key X exists



Does key  
X exist



Memory



**true positive**


Data



**If key X does not exist**




Does key  
X exist



**No**


Memory




Data




# If key X does not exist



If key X does not exist




If key X does not exist



$\epsilon$  - false positive rate - FPR


# Bloom Filters

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



# ***k* hash functions**


0 0 0 0 0 0 0 0 0 0



**bitmap**

# Inserts?


0 0 0 0 0 0 0 0 0 0






Inserts: **hash key to k random bits**

**insert(X)**




0 0 0 0 0 0 0 0 0 0




Inserts: hash key to k random bits

**Set from 0 to 1 or keep 1**

insert(X)




0 0 0 0 **1** 0 0 0 **1** 0




Inserts: hash key to k random bits

**Set from 0 to 1 or keep 1**

**insert(Y)**




0 0 **1** 0 **1** 0 0 0 **1** 0




# Queries?

0 0 **1** 0 **1** 0 0 0 **1** 0



Queries: **return positive if all hashed bits are 1s**

0 0 **1** 0 **1** 0 0 0 **1** 0




Queries: return positive if all hashed bits are 1s


**check(X)**



**True  
positive**



0 0 **1** 0 **1** 0 0 0 **1** 0




Queries: return positive if all hashed bits are 1s


**check(Z)**



**True  
negative**



0 0 **1** 0 **1** 0 0 0 **1** 0




Queries: return positive if all hashed bits are 1s


**check(Q)**



**False  
Positive**



0 0 **1** 0 **1** 0 0 0 **1** 0






# Can the filter ever return a false negative?

i.e., report that an existing key does not exist...


0 0 **1** 0 **1** 0 0 0 **1** 0



Can the filter ever return a false negative?

**No - the corresponding bits will always be 1s**

0 0 **1** 0 **1** 0 0 0 **1** 0




Can the filter ever return a false negative?

No - the corresponding bits will always be 1s


**False negatives are not allowed...**

0 0 **1** 0 **1** 0 0 0 **1** 0



# How many hash functions should we use?

0 0 0 0 0 0 0 0 0 0




# How many hash functions should we use?

$h_1$




0 1 0 0 0 0 0 0 0 0



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

# How many hash functions should we use?




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


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

# How many hash functions should we use?

$h_1$                       ...                       $h_x$



1 1 1 1 1 1 1 1 1 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).

**But too many hash functions wind up increasing the FPR.**


# How many hash functions should we use?



(Drawn for a filter using 10 bits per entry)



# How many hash functions should we use?




**Optimal # hash functions =  $\ln(2) \cdot M$**

( $M$  is the number of bits per entry)


assuming the optimal # hash functions,

$$\text{false positive rate} = 2^{-M \cdot \ln(2)}$$




**deletes?**

0 0 **1** 0 **1** 0 0 0 **1** 0




deletes?




# deletes?

Delete Y X



0 0 **0** 0 **1** 0 0 0 **0** 0




deletes?

Query X



0 0 0 0 **1** 0 0 0 **0** 0




deletes?

Query X false negative



0 0 0 0 **1** 0 0 0 **0** 0



# Research Challenges

**deletes?**






# Research Challenges

deletes?



**Expansions?**




# Research Challenges

deletes?



Expansions?




**Range?**



deletes?



Expansions?



Range?



---


**All at the same time?**

**Subject to:**

deletes?



Expansions?



Range?



Subject to:

**High  
Performance**



**Small  
Size**




**Good  
Accuracy**



# Orca Lab





**CSC443: Database  
System Technology**



**CSC2525: Research topics  
in Database Management**