Introduction to Randomized Algorithms IV

Joaquim Madeira

Version 0.0 – December 2019

Overview

- Set membership queries
- Hash Tables – A quick review
- Hash Functions – A quick review
- Bloom Filters
- Counting Bloom Filters
- More recent AMQ Filters

Set Membership

- Given an arbitrary sized string s and a set S
- Does s belong to S ?

- Easy answer for small sets !
  - Complexity ?

- BUT “difficult” answer for huge sets !
  - E.g., Big-Data applications

- Approximate Membership Queries (AMQ)

Hash Tables

- Data structure for storing key-value pairs
- No ordering !!
- BUT, fast access !!
- No duplicate keys !!

Hash Tables

- Two main operations :
- Insert (put) a key-value pair into the table
  - If key already exists, update the value
- Search for (get) the value associated with a given key

Hash Tables

- Additional operations :

- contains(key)
- delete(key)
- is_empty()
- Keys iterator
- …

Hashing

- To reference key-value pairs stored in a table
- Perform arithmetic operations that transform search keys into array indices
  - FAST !!
- Ideally, different keys would map to different indices
- BUT, collisions do occur !!

HashTables – Time complexity

- The time complexity of searches by hashing can be as low as O(1) or as high as O(N)

- Worst-case ?
- Distinct keys $K_i \neq Kj$ collide: $h(K_i) = h(K_j)$
- The entire table must be searched to find the correct entry
- Or to conclude it is not there !

Hash Functions

- Pseudo-random mathematical functions used to compute indices for table look-up
  - Keys are mapped to small integers
- Indices should be evenly distributed
  - Even if there are regularities in the data

- There are many hash functions
  - With different degrees of complexity
  - And with differences in performance
  - For different applications

Simple Hash Functions

- Division method
  - Choose a prime m that isn't close to a power of 2
  - $h(k) = k \bmod m$
  - Works badly for many types of patterns in the input data
- Knuth's variant
  - $h(k) = k(k+3) \bmod m$
  - Supposedly works much better than the raw division method

Simple Hash Functions

```
def hash(astring, tablesize):
    sum = 0
    for pos in range(len(astring)):
        sum = sum + ord(astring[pos])

    return sum%tablesize
```

- Anagrams will be given the same values…

Hash Functions – DJB31MA

```
uint hash(const uchar* s, int len, uint seed)
{
    uint h = seed;
    for (int i=0; i < len; ++i)
        h = 31 * h + s[i];
    return h;
}
```

Non-cryptographic Hash Functions

- Suitable for hash table lookup but not for crytography / secure uses
- Fast computation

- FNV – Fowler-Noll-Vo hash function
- Murmur Hash
  - Multiply and rotate
- …

Universal Hashing

- Issue
  - There always exist keys that are mapped to the same integer / index

- Consider a set of hash functions H
- H is universal (good), if
  - For all keys $0 \leq i \leq j < M$
  - Probability $(h(i) = h(j)) \leq 1 / M$, for h randomly selected from H

Approximate Membership Queries

- Given a set $S = \{x_1, x_2, \ldots, xn\}$
- Answer queries of the form: $Is\ y\ in\ S\ ?$

- Data structure should be FAST and SMALL
  - Faster than searching through S
  - Smaller than explicit representation

Approximate Membership Queries

- How to get speed and size improvements ?
- Allow some probability of error !!
- False positives
  - $y \notin S$ but reporting $y \in S$
- False negatives
  - $y \in S$ but reporting $y \notin S$

Bloom Filters

- B. H. Bloom, 1970
- Use hash functions to determine approximate set membership
- Allow for fast set membership tests on very large data sets

- Applications
  - Spell-Checking / Text Analysis
  - Network monitoring
  - …

Application – Spell-Checkers

- Determine if candidate words are members of the set of words in a dictionary
- The Bloom filter should be large enough to allow the inclusion of additional words by the user

Application – Email Spam

- We know 1 billion “good” email addresses
- If an email comes from one of these, it is **NOT** spam
- How check for spam in a FAST way ?

Application – Web-Caching

- Bloom filters are used in WWW caching proxy servers
- Proxy servers intercept requests from clients and either fulfill the requests themselves or re-issue them to servers

Application – Text Analysis

- Find related passages in different reports
- Constructing a Bloom Filter of all the words in each passage
- Computing the normalized dot product of all Bloom filter pairs
- The result of every dot product is a similarity measure

Bloom Filters

- ***Is y in S*** ?

- A Bloom filter
  - Provides an answer in constant time
    - Time to hash
- Uses a small amount of memory space
- BUT, with some small probability of being wrong !

1st – Register the elements of set S

Start with an m bit array, filled with 0s.

B	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Hash each item x_j in S k times. If $H_i(x_j) = a$, set $B[a] = 1$.

B	0	1	0	0	1	0	1	0	0	1	1	1	0	1	1	0

[Mitzenmacher]

2nd – Process the queries

To check if y is in S, check B at $H_i(y)$. All k values must be 1.

B	0	1	0	0	1	0	1	0	0	1	1	1	0	1	1	0

Possible to have a false positive; all k values are 1, but y is not in S.

B	0	1	0	0	1	0	1	0	0	1	1	1	0	1	1	0

[Mitzenmacher]

Basic operations

- Initialization
  - Clear all cells
- Insertion
  - Compute the values of k hash functions
  - Set the corresponding cells, if needed
  - It takes constant time, but proportional to k

Basic operations

- Membership test
  - Compute the values of k hash functions
  - Check if the corresponding cells have been set
  - If any such cell is not set, the searched element is not a member of the set

- Worst-case ?
- Checking all k cells !
  - Set elements and false positives

Bloom Filter – Simple Demos

- Bloom Filters by Example
  - http://billmill.org/bloomfilter-tutorial/
- Bloom Filters
  - https://www.jasondavies.com/bloomfilter/

Bloom Filters – Behaviour

- Deterministic hash functions !
- No attempt to solve hashing collisions !
- Can we get false negatives ?
- Probability of false positives ?
- How to minimize ?

Bloom Filter – Parameters

- The behaviour of a Bloom filter is determined by four parameters

- n set elements registered in B
- $m = c \times n$ cells in B (i.e., bits)
- k independent, random hash functions
- f is the fraction of cells set to 1

Bloom Filter – Parameters

- How to choose m, the size of the filter ?
- How to choose k, the number of hash functions ?
- How do we choose the best k value ?

Probabilities – After 1 insertion

- Initially all bits are set to zero
- Inserting one element

- What is the probability of $b_i = 1$, after using the first hash function ?
  - Equal probability for any cell

$$P(b_i = 1) = \frac{1}{m}$$

$$P(b_i = 0) = 1 - \frac{1}{m}$$

Probabilities – After 1 insertion

- After computing the k hash functions and setting k cells

$$P(b_i = 0) = \left(1 - \frac{1}{m}\right)^k$$

Probabilities – After n insertions

- After inserting all n set elements, by computing each time k hash values
  - Assuming independence

$$P(b_i = 0) = \left(1 - \frac{1}{m}\right)^{k \times n}$$

Probabilities – After n insertions

$$P(b_i = 0) = \left(1 - \frac{1}{m}\right)^{k \times n}$$

$$P(b_i = 1) = 1 - a^k, \qquad a = \left(1 - \frac{1}{m}\right)^{n}$$

Probability of a false positive

- Testing the membership of an item not in S entails a positive answer
  - Corresponding k bits are set to 1

- The probability of that happening is

$$p = \left(1 - a^k\right)^k$$

$$p \approx \left(1 - e^{-kn/m}\right)^k$$

Example

- n = 1 billion items, m = 8 billion bits
- k = 1 : $p \approx \left(1 - e^{-1/8}\right) = 0.1175$
- k = 2 : $p \approx \left(1 - e^{-2/8}\right)^2 = 0.0493$
- What happens as we keep increasing k ?

Optimal value of k

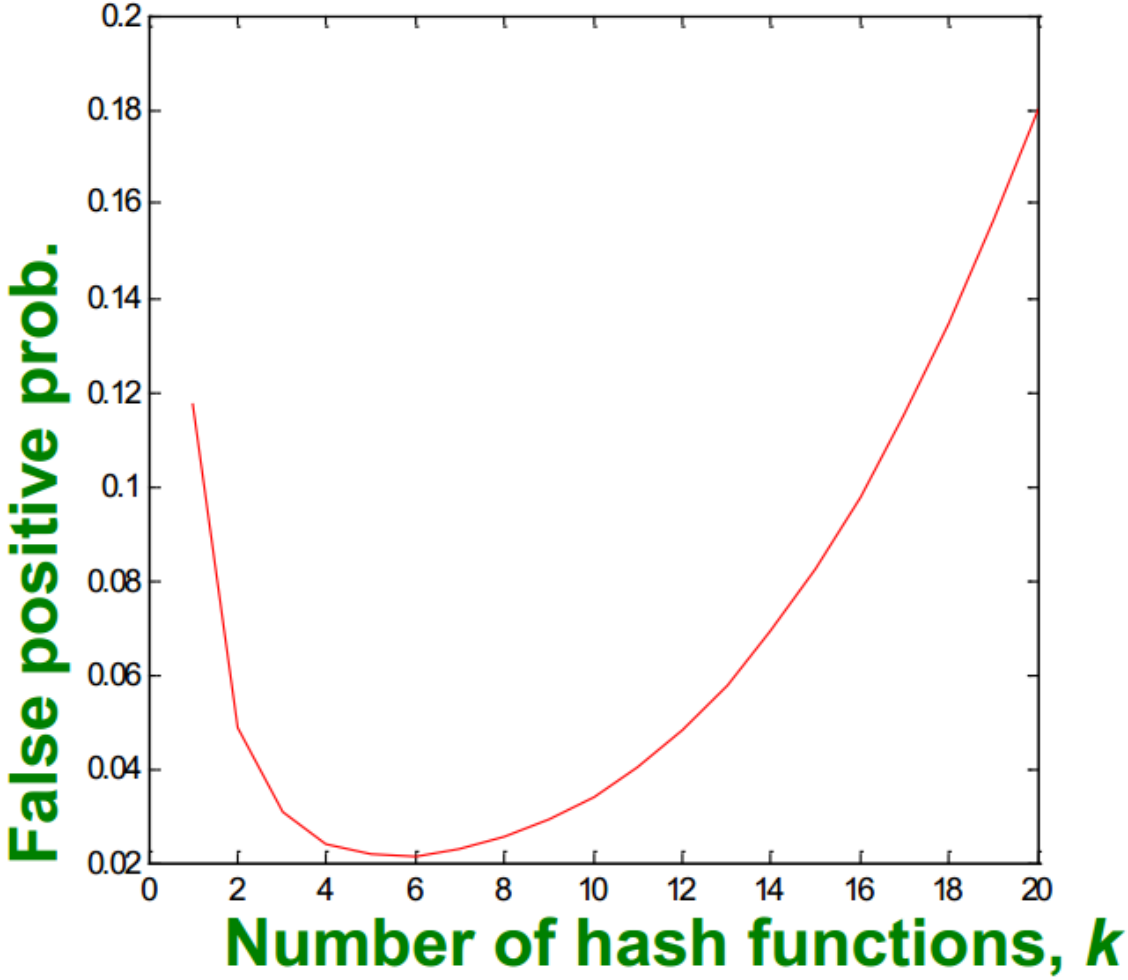

Optimal value of k

- To determine the value of k that minimizes p we minimize log p, which is more tractable
- And get

$$k_{opt} \approx \frac{m}{n} \times \ln 2 \approx 0.693 \times \frac{m}{n}$$

- Use the closest integer to k_{opt}
- For the previous example : $k_{opt} \approx 5{,}54 \approx 6$

Which Hash Functions ?

- No need to use cryptographic hash functions !
- You can simulate k hash functions by simply combining two hash functions
  - Kirsch and Mitzenmacher (2006)
- Compute one base hash function on unsigned 64-bit numbers
- Take the upper half and the lower half of that value and return them as *two* 32 bit numbers

Bloom Filters – Wrap-up

- No false negatives and limited memory usage
  - Great for pre-processing before more expensive checks
- Suitable for hardware implementation
  - Hash computations can be parallelized
- Error rate can be decreased by increasing the number of hash functions and allocated memory space

Bloom Filters – Wrap-up

- Useful for applications where an imperfect set membership test can be helpfully applied to a large data set of unknown composition
- Advantage over hash tables is Bloom filter speed and error rate

Bloom Filters – Pending Issues

- Cannot represent multi-sets
  - I.e., sets with repeated elements
- Cannot query the multiplicity of an item
- Deleting an item is not possible !

Counting Bloom Filters

- Multi-set representation
- Now, each filter cell is a w-bit counter
  - w = 4 seems to be enough for most applications

Counting Bloom Filters

- To insert an element, increase the value of each corresponding cell
- Test membership checks if each of the required cells is non-zero

Counting Bloom Filters

- To delete an element, decrease the value of each corresponding cell
- Deletions necessarily introduce false negative errors !!
  - How ?

Counting Bloom Filters

- To retrieve the count of an element :
- Compute its set of counters
- And return the minimum value as a frequency estimate

Counting Bloom Filters

Start with an m bit array, filled with 0s.

B	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Hash each item x_j in S k times. If $H_i(x_j) = a$, add 1 to $B[a]$.

B	0	3	0	0	1	0	2	0	0	3	2	1	0	2	1	0

[Mitzenmacher]

Counting Bloom Filters

To delete x_j decrement the corresponding counters.

B	0	2	0	0	0	0	2	0	0	3	2	1	0	1	1	0

Can obtain a corresponding Bloom filter by reducing to 0/1.

B	0	1	0	0	1	0	1	0	0	1	1	1	0	0	1	0

[Mitzenmacher]

Counting Bloom Filters – Issues

- Counter overflow
  - No more increments after reaching $2^w - 1$
  - BUT, now we have undercounts !!
- Choice of counter width w
  - A large w diminishes space savings and introduces unused space (many zeros)
  - A small w quickly leads to maximum values
  - Trade-off…

Counting Bloomm Filters in Practice

- If insertions/deletions are rare compared to look-ups
  - Keep a CBF in “off-chip memory”
  - Keep a BF in “on-chip memory”
  - Update the BF when the CBF changes
- Keep space savings of a Bloom filter
- But can deal with deletions
- Popular design for network devices

References

- J. Leskovec, A. Rajaraman and J. D. Ullman, *Mining of Massive Datasets*, 2014 – Chapter 4
- B. H. Bloom, Space/Time Trade-offs in Hash Coding with Allowable Errors, *Commun. ACM*, July 1970
- J. Blustein and A. El-Maazaw, Bloom Filters – A Tutorial, Analysis, and Survey, TR CS 2002-10, Dalhousie University, Halifax, NS, Canada, December 2002
- A. Broder and M. Mitzenmacher, Network Applications of Bloom Filters: A Survey, *Internet Mathematics*, Vol. 1, N. 4, 2004

Bloom Filters – Main limitation

- What happens if the filter is too big to fit in main memory ?
- Store some parts of the filter in HDD or SSD
- BUT, random accesses are slower in disk…
- Significantly bad performance !

Bloom Filters – Main limitation

- IDEAS:
  - Random access only once
  - Store each element's data really close
- BUT
  - When using just 1 hash function
  - Have to deal with high collision probability
- HOW ? --- The Quotient Filter

The Quotient Filter

- 2011 : Michael Bender et al.
- Space-efficient probabilistic data structure for AMQ
- Implements a set with 4 operations
  - Add element
  - Delete element
  - Test whether an element is a member
  - Test whether an element is not a member

The Quotient Filter

- BUT, do not actually store each element
- Just store a p-bit fingerprint for each element
- Using just one hash function
- Compact open hash table with $\boldsymbol{m = 2^q}$ buckets

Quotienting

- The fingerprint f (p bits) is partitioned
- Division by 2^r, i.e., shift to the right
  - The remainder f_r : r least significant bits
  - The quotient f_q : q most significant bits
- The quotient indexes a table bucket
- The remainder is stored in that bucket

Soft-collision handling

- Collisions do occur !!
- Different fingerprints have the same quotient :

$$f_q = f^*_q$$

- All remainders of fingerprints with the same quotient are stored contiguously in a run
- If necessary, a remainder is shifted forward from its original location
  - Stored in a subsequent bucket
  - Wrapping around at the end of the table

3 auxiliary bits per bucket

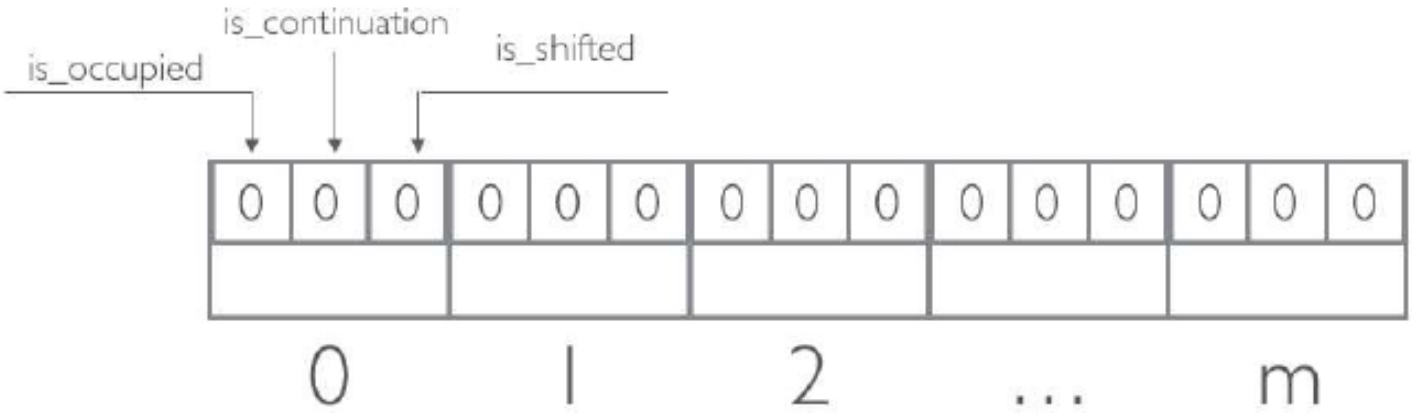

[Gakhov]

- Initially set to zero
- is_occupied : $f_q = j$ for some stored fingerprint
  - j is the canonical bucket
- is_continuation : occupied, but not by first remainder in a run
- is_shifted : remainder in the bucket is not in its canonical bucket

Membership testing

- Given the searched element
- Use hash function to compute its fingerprint
- Compute quotient f_q and remainder f_r
- If bucket f_q is not occupied : item definitely not in the filter !!

Membership testing

- If bucket f_q is occupied
  - Scan left to locate first bucket with is_shifted = 0
  - Scan right to identify the quotient's run
  - For each bucket in the run, compare the stored remainder with f_r
  - If found, the searched element is (probably) in the filter
  - Else, it is definitely not in the filter

Adding an element

- Given the searched element
- Use hash function to compute its fingerprint
- Compute quotient f_q and remainder f_r
- Follow a path similar to the testing procedure, until certain that element not in the filter

Adding an element

- Choose bucket in the current run and insert the remainder f_r
  - Keep the sorted order !!
- Shift forward all remainders at or after the chosen bucket
  - Update the buckets' bits

Example – Adding elements

- Quotient size ***q = 3***
- 32-bit signed MurmurHash3
- Add
  - $f_q(\text{amsterdam}) = 1,\ f_r(\text{amsterdam}) = 164894540$
  - $f_q(\text{berlin}) = 4,\ f_r(\text{berlin}) = -89622902$
  - $f_q(\text{london}) = 7,\ f_r(\text{london}) = 232552816$

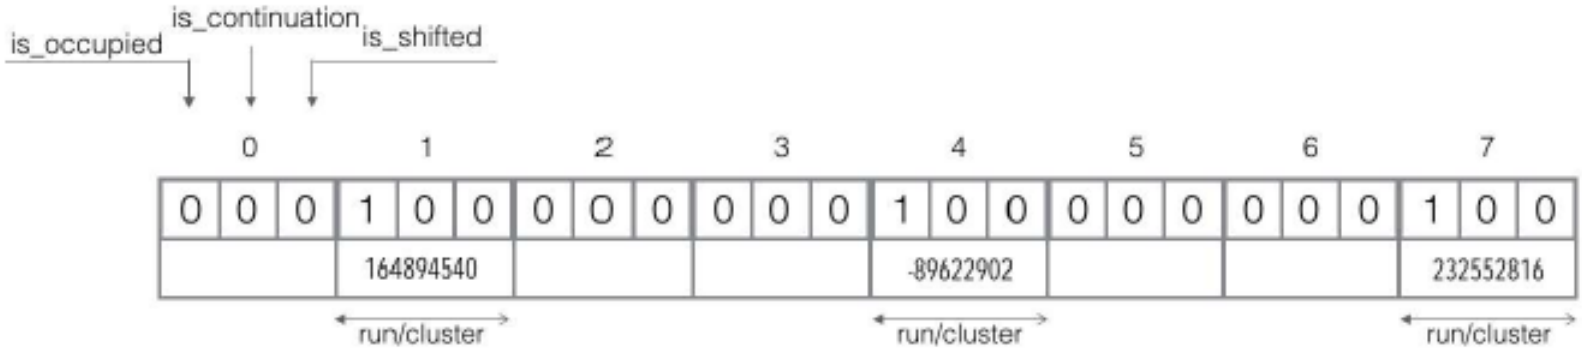

[Gakhov]

Example – Adding elements

- Add

$$f_q(\text{madrid}) = 1,\ f_r(\text{madrid}) = 249059682.$$

$$f_q(\text{ankara}) = 2,\ f_r(\text{ankara}) = 62147742.$$

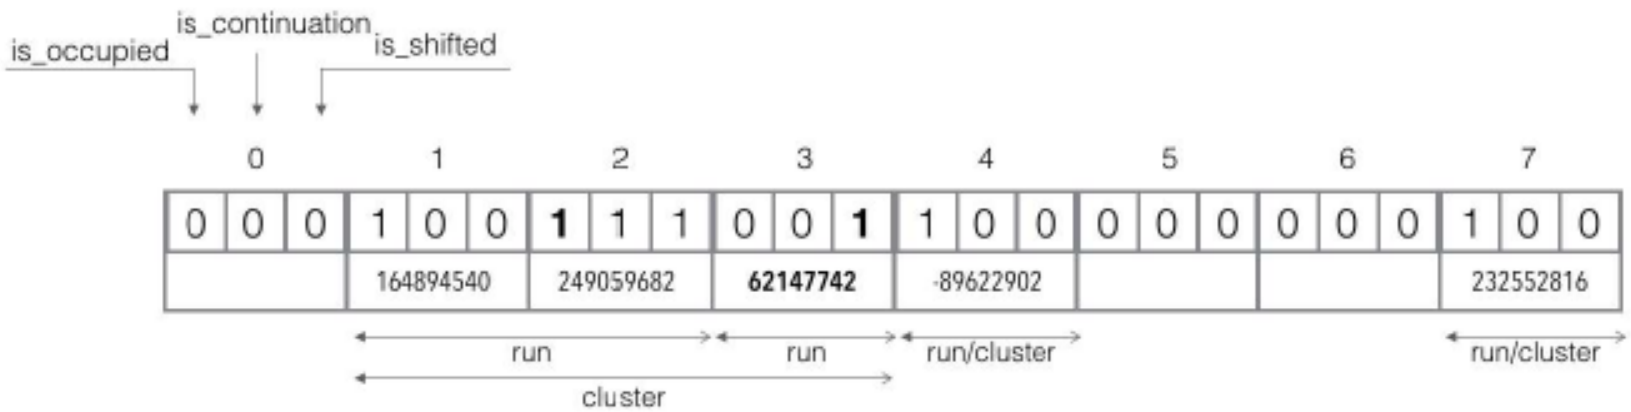

[Gakhov]

Example – Testing membership

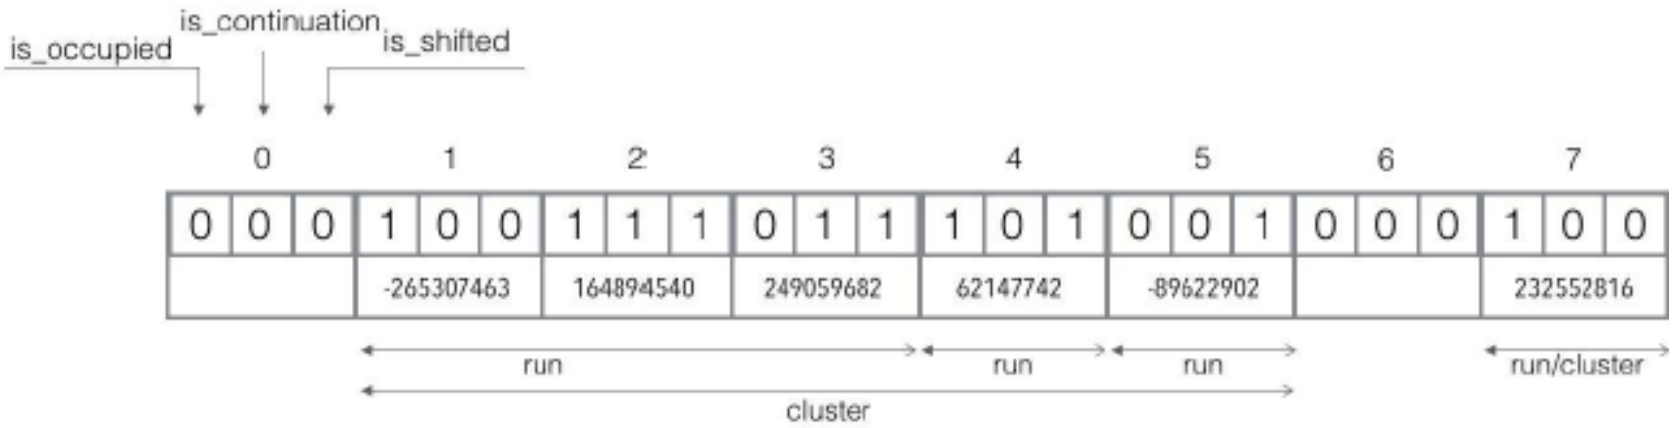

$$f_q(\text{ferret}) = 1,\ f_r(\text{ferret}) = 122150710.$$

$$f_q(\text{berlin}) = 4,\ f_r(\text{berlin}) = -89622902.$$

[Gakhov]

Quotient Filter - Features

- False positives are possible, but with low probability

$$P(e \in S \mid e \notin S) \leq \frac{1}{2^r}$$

- False positive probability can be tuned
- False negatives are not possible

Quotient Filter - Features

- Hash function should generate uniformly distributed fingerprints
- The length of most runs is O(1)
- Probable that most runs have length O(log m)
- Efficient filter for large number of elements

Quotient Filter vs Bloom Filter

- QFs are about 20% bigger
- QFs are faster : evaluate single hash function
- Comparison results (M. Bender)
  - **inserts**. BF: 690 000 inserts per second, QF: 2 400 000 insert per second
  - **lookups**. BF: 1 900 000 lookups per second, QF: 2 000 000 lookups per second
- QFs support deletion !!

2014 – Cuckoo Filters

Cuckoo Filter: Practically Better Than Bloom

Bin Fan, David G. Andersen, Michael Kaminsky[†], Michael D. Mitzenmacher[‡]
Carnegie Mellon University, [†]Intel Labs, [‡]Harvard University

ABSTRACT

In many networking systems, Bloom filters are used for high-speed set membership tests. They permit a small fraction of false positive answers with very good space efficiency. However, they do not permit deletion of items from the set, and previous attempts to extend "standard" Bloom filters to support deletion all degrade either space or performance.

We propose a new data structure called the *cuckoo filter* that can replace Bloom filters for approximate set membership tests. Cuckoo filters support adding and removing items dynamically while achieving even higher performance than Bloom filters. For applications that store many items and

2017 – Counting Quotient Filters

A General-Purpose Counting Filter: Making Every Bit Count

Prashant Pandey, Michael A. Bender, Rob Johnson, and Rob Patro
Stony Brook University

ABSTRACT

Approximate Membership Query (AMQ) data structures, such as the Bloom filter, quotient filter, and cuckoo filter, have found numerous applications in databases, storage systems, networks, computational biology, and other domains. However, many applications must work around limitations in the capabilities or performance of current AMQs, making these applications more complex and less performant. For example, many current AMQs cannot delete or count the number of occurrences of each input item, take up large amounts of space, are slow, cannot be resized or merged, or have poor locality of reference and hence perform poorly when stored on SSD or disk.

This paper proposes a new general-purpose AMQ, the counting quotient filter (CQF). The CQF supports approximate membership testing and counting the occurrences of items in a data set. This

2018 – Morton Filters

Morton Filters: Faster, Space-Efficient Cuckoo Filters via Biasing, Compression, and Decoupled Logical Sparsity

Alex D. Breslow
Advanced Micro Devices, Inc.
AMD Research
2485 Augustine Drive

Nuwan S. Jayasena
Advanced Micro Devices, Inc.
AMD Research
2485 Augustine Drive

ABSTRACT

Approximate set membership data structures (ASMDSs) are ubiquitous in computing. They trade a tunable, often small, error rate (ϵ) for large space savings. The canonical ASMDS is the Bloom filter, which supports lookups and insertions but not deletions in its simplest form. Cuckoo filters (CFs), a recently proposed class of ASMDSs, add deletion support and often use fewer bits per item for equal ϵ.

This work introduces the Morton filter (MF), a novel ASMDS that introduces several key improvements to CFs. Like CFs, MFs support lookups, insertions, and deletions, but improve their respective throughputs by 1.3× to 2.5×, 0.9× to 15.5×, and 1.3× to 1.6×. MFs achieve these improve-

References

- A. Gakhov. Probabilistic data structures. Quotient filter
- Michael A. Bender et al. Don't Thrash: How to Cache your Hash on Flash. 3rd USENIX Workshop HotStorage'11, June 14-17, 2011
- Michael A. Bender et al. Don't Thrash: How to Cache Your Hash on Flash. In Proceedings of the VLDB Endowment (2012), Vol. 5 (11). Pp. 1627-1637

Acknowledgments

- An earlier version of some of these slides was developed by Professor Carlos Bastos

- Part of the slides adapted from original slides of
  - J. Leskovec, A Rajaraman and J. Ullman – Mining of Massive Datasets – www.mmds.org
  - M. Mitzenmacher, Bloom Filters and Such – 2014 Summer School on Hashing, Copenhagen, DK