Cuckoo Filter: Practically Better Than Bloom

Bin Fan (CMU/Google)
David Andersen (CMU)
Michael Kaminsky (Intel Labs)
Michael Mitzenmacher (Harvard)

Carnegie Mellon
Parallel Data Laboratory

What is Bloom Filter? A Compact Data Structure Storing Set-membership

- Bloom Filters answer "is item x in set Y" by:
  - "definitely no", or
  - "probably yes" with probability ε to be wrong

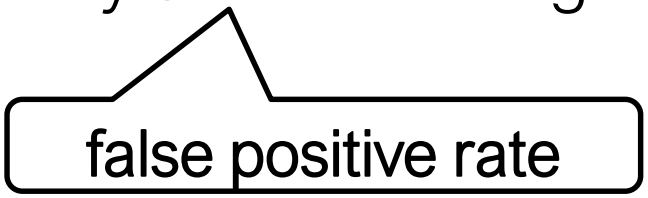

- Benefit: not always precise but highly compact
  - Typically a few bits per item
  - Achieving lower ε (more accurate) requires spending more bits per item

Example Use: Safe Browsing

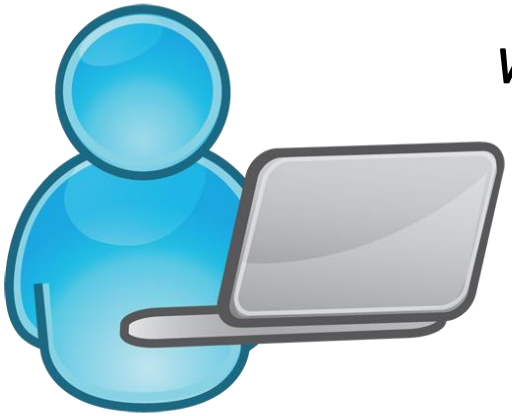

www.binfan.com

Example Use: Safe Browsing

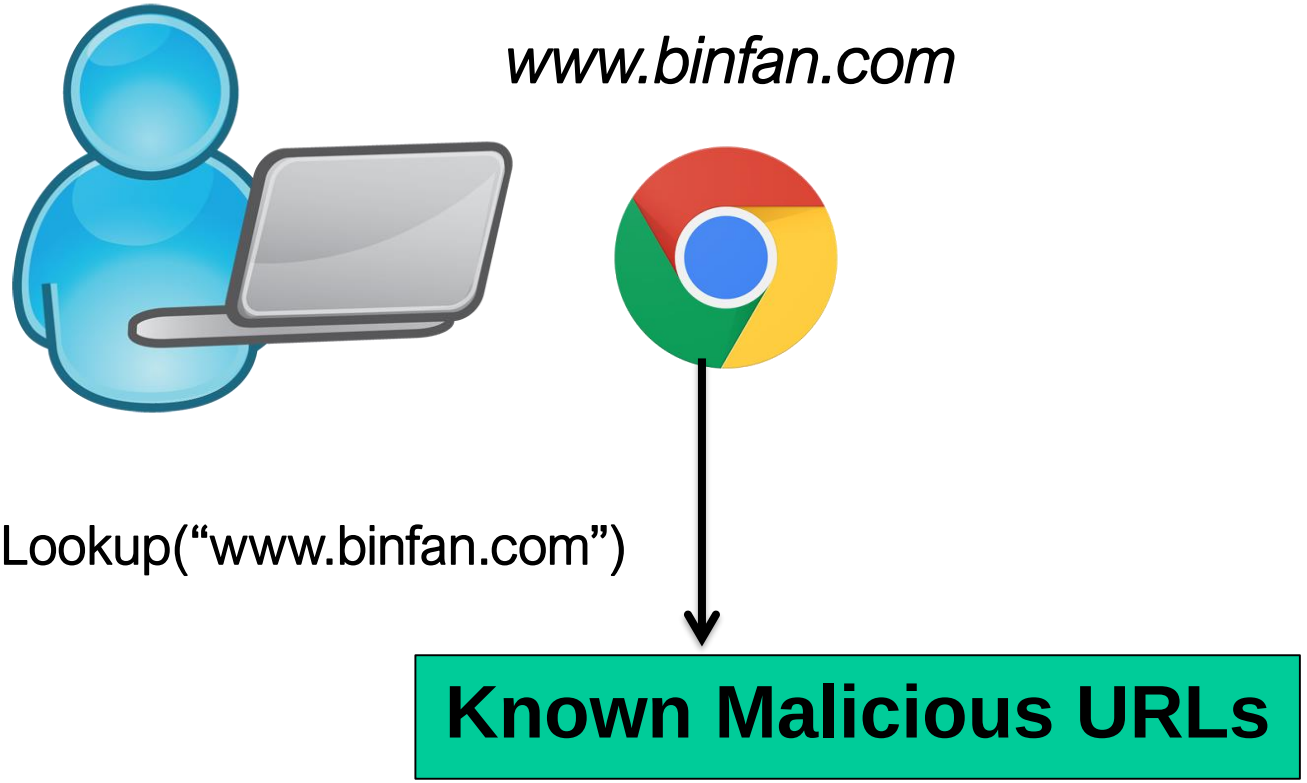

Example Use: Safe Browsing

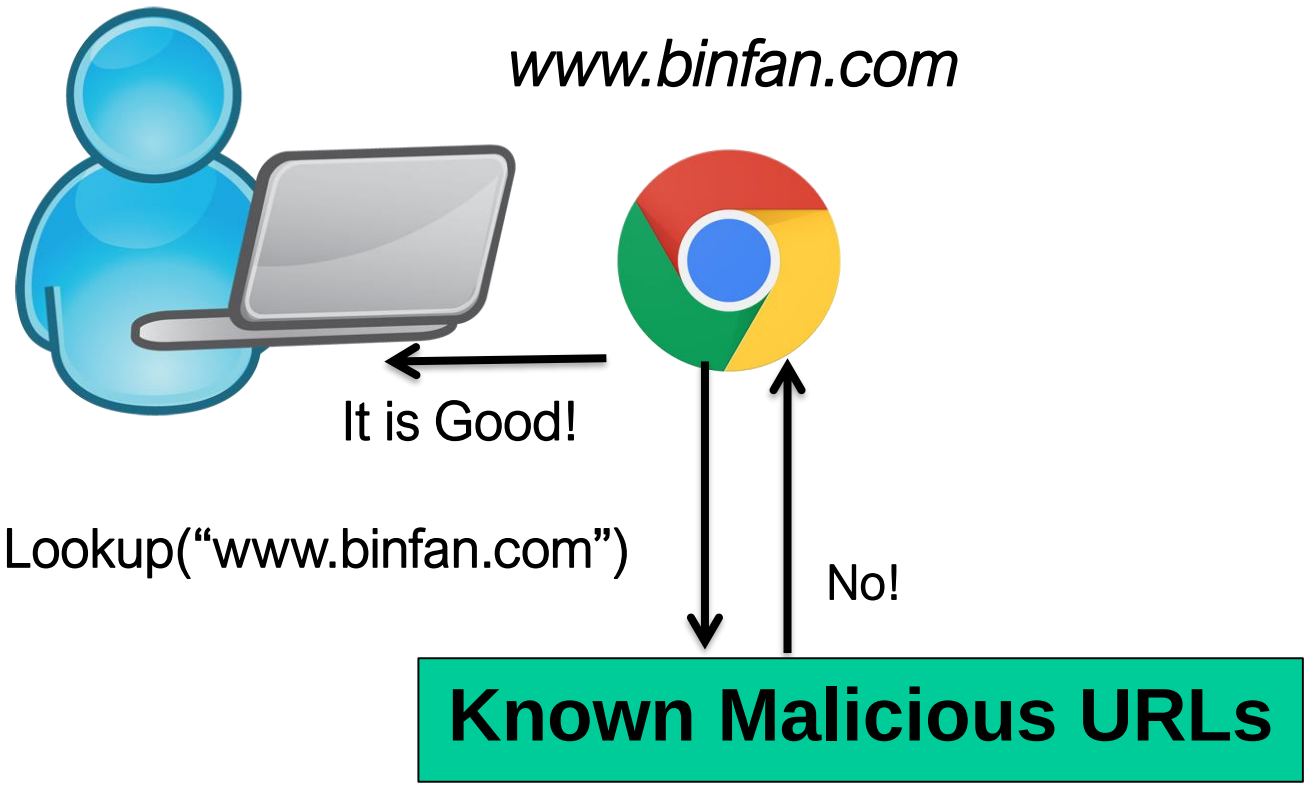

Example Use: Safe Browsing

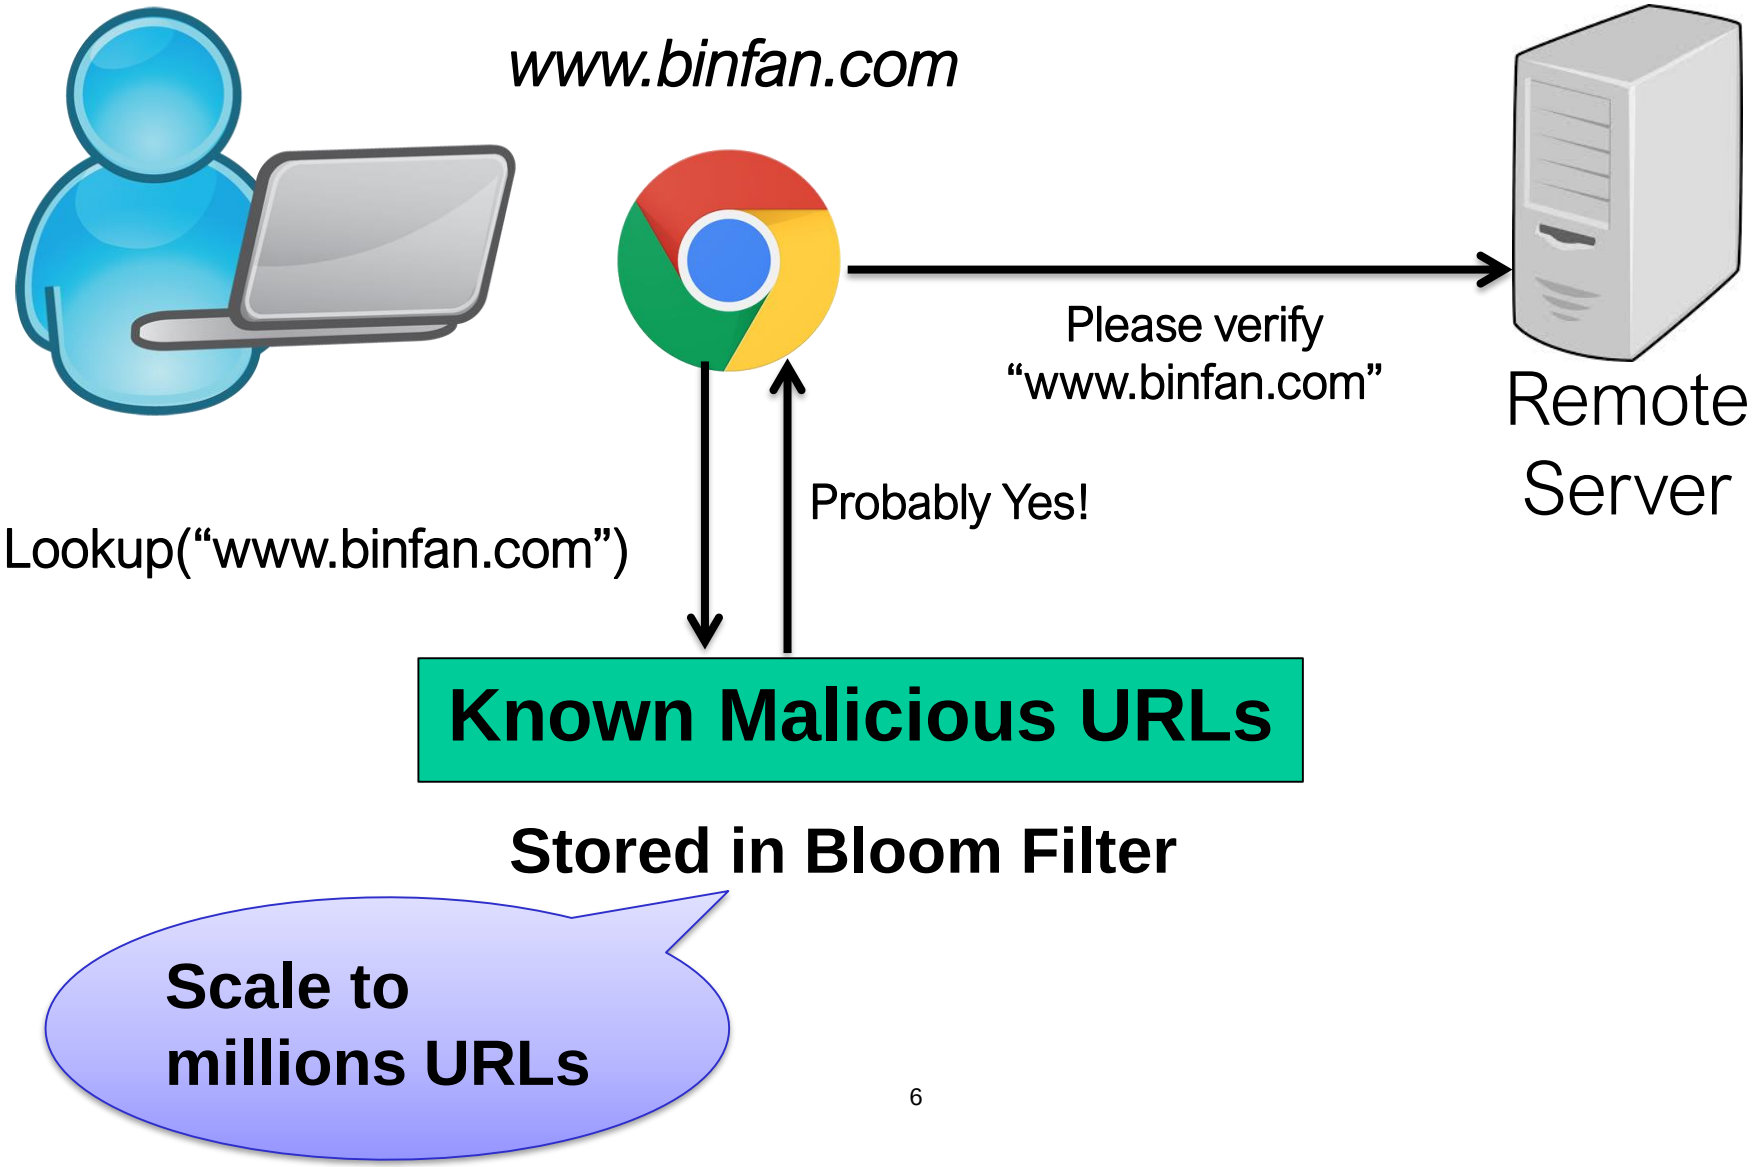

Bloom Filter Basics

A Bloom Filter consists of ***m*** bits and ***k*** hash functions

Example: ***m*** = 10, ***k*** = 3

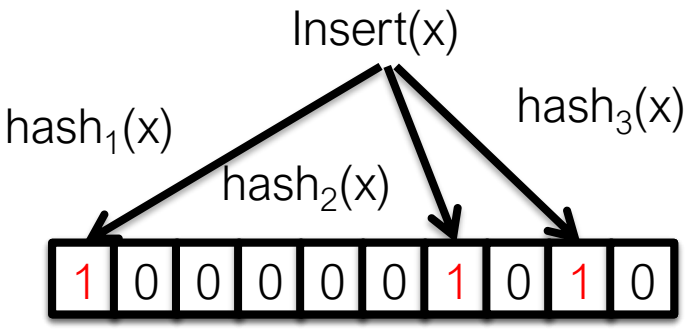

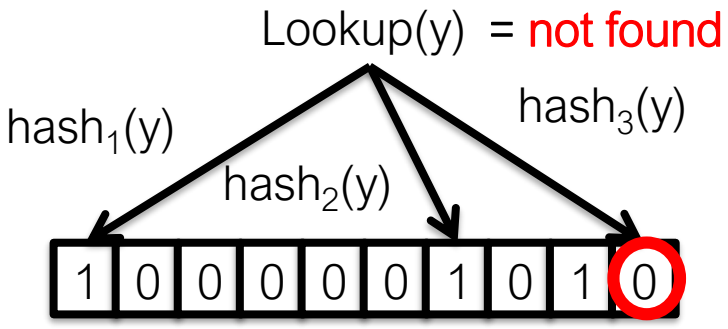

Succinct Data Structures for Approximate Set-membership Tests

	High Performance	Low Space Cost	Delete Support
Bloom Filter	✓	✓	✗
Counting Bloom Filter	✓	✗	✓
Quotient Filter	✗	✓	✓

Can we achieve all three in practice?

Outline

- Background
- Cuckoo filter algorithm
- Performance evaluation
- Summary

Basic Idea: Store Fingerprints in Hash Table

- **Fingerprint**(x): A hash value of x
  - Lower false positive rate ε, longer fingerprint

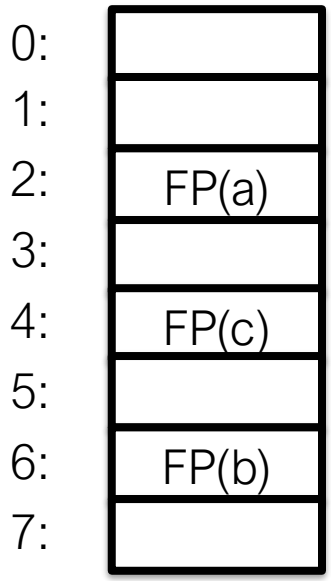

Basic Idea: Store Fingerprints in Hash Table

- **Fingerprint**(x): A hash value of x
  - Lower false positive rate ε, longer fingerprint
- **Insert**(x):
  - add **Fingerprint(x)** to hash table

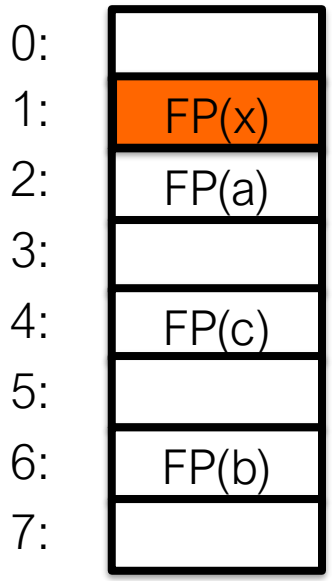

Basic Idea: Store Fingerprints in Hash Table

- **Fingerprint**(x): A hash value of x
  - Lower false positive rate ε, longer fingerprint
- **Insert**(x):
  - add **Fingerprint(x)** to hash table
- **Lookup**(x):
  - search **Fingerprint**(x) in hashtable

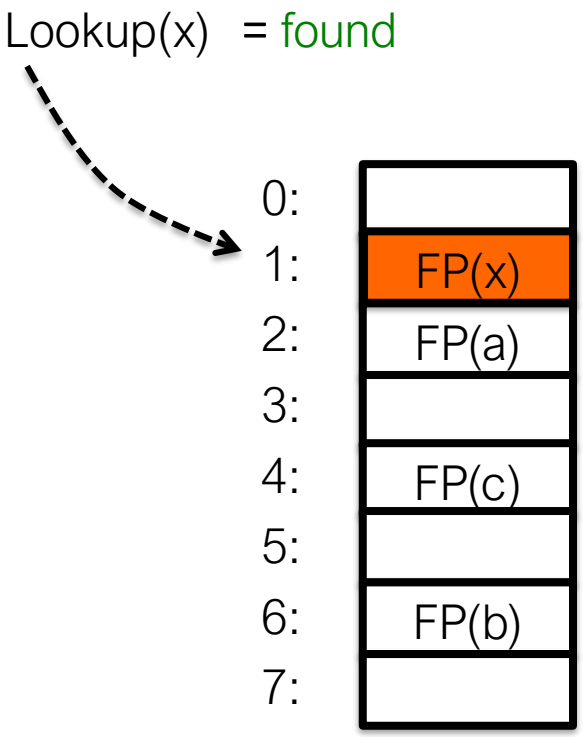

Basic Idea: Store Fingerprints in Hash Table

- **Fingerprint**(x): A hash value of x
  - Lower false positive rate ε, longer fingerprint
- **Insert**(x):
  - add **Fingerprint(x)** to hash table
- **Lookup**(x):
  - search **Fingerprint**(x) in hashtable
- **Delete**(x):
  - remove **Fingerprint**(x) from hashtable

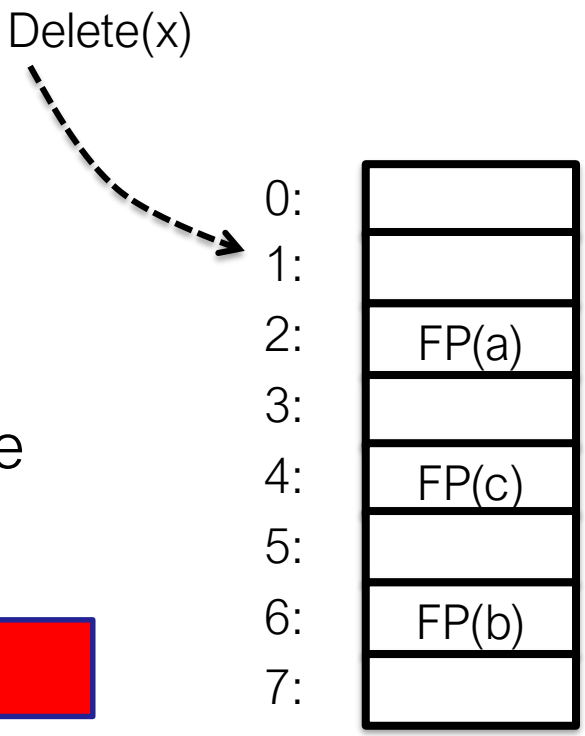

How to Construct Hashtable?

Strawman

(Minimal) Perfect Hashing: No Collision but Update is Expensive

- Perfect hashing: maps all items with no collisions

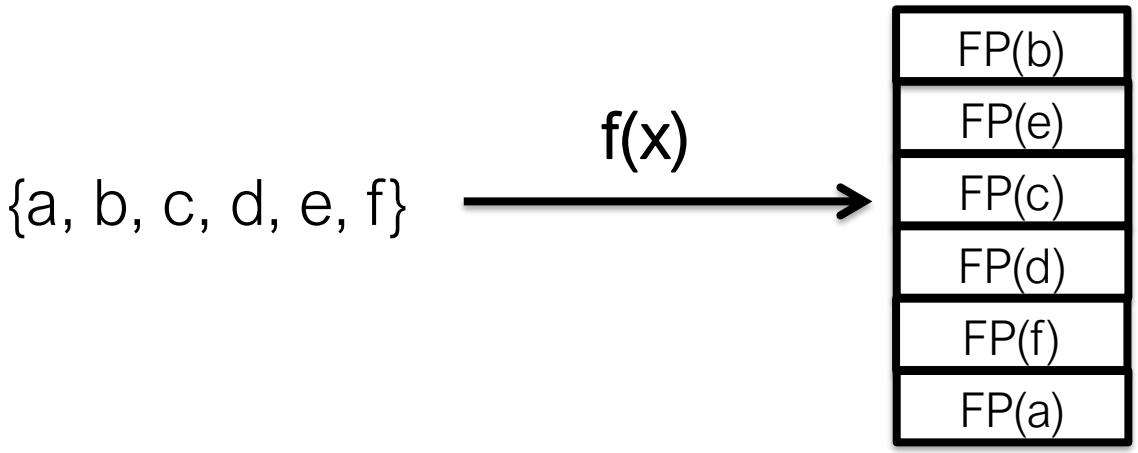

Strawman

(Minimum) Perfect Hashing: No Collision but Update is Expensive

- Perfect hashing: maps all items with no collisions

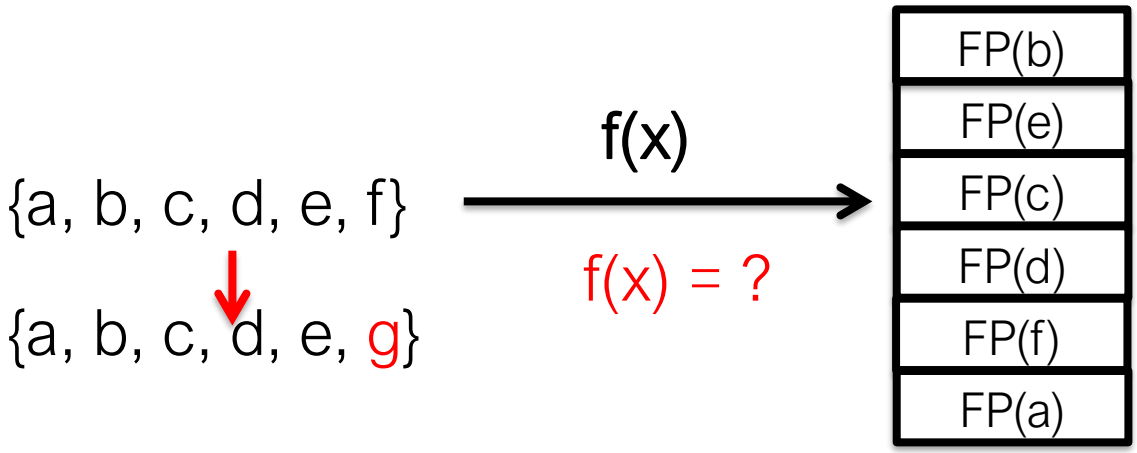

- Changing set must recalculate f ➔ high cost/bad performance of update

Strawman

Convention Hash Table: High Space Cost

- Chaining :

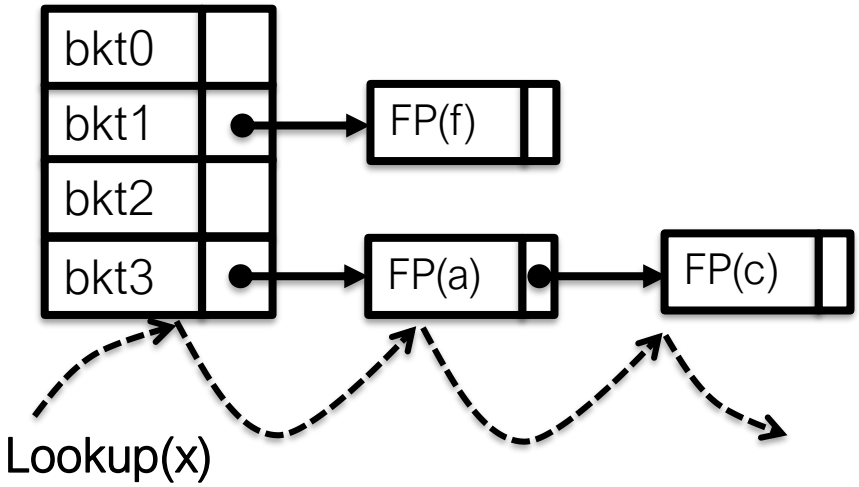

- Pointers ➔
  low space utilization

- Linear Probing

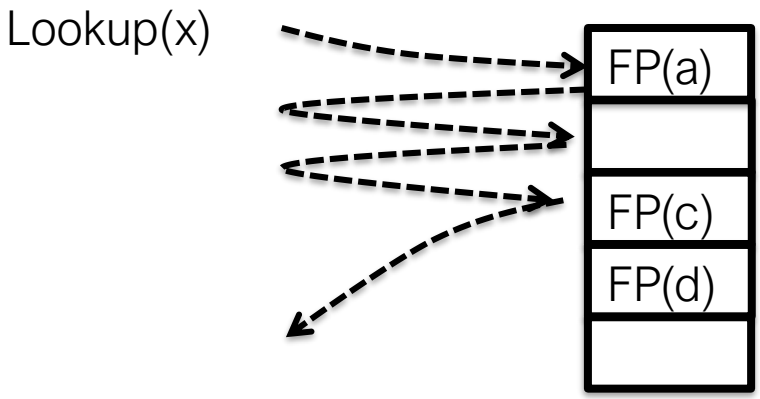

- Making lookups O(1) requires large % table empty ➔
  low space utilization
- Compare multiple fingerprints sequentially ➔
  more false positives

Cuckoo Hashing[Pagh2004] Good But ..

- High Space Utilization
  - 4-way set-associative table: >95% entries occupied
- Fast Lookup: O(1)

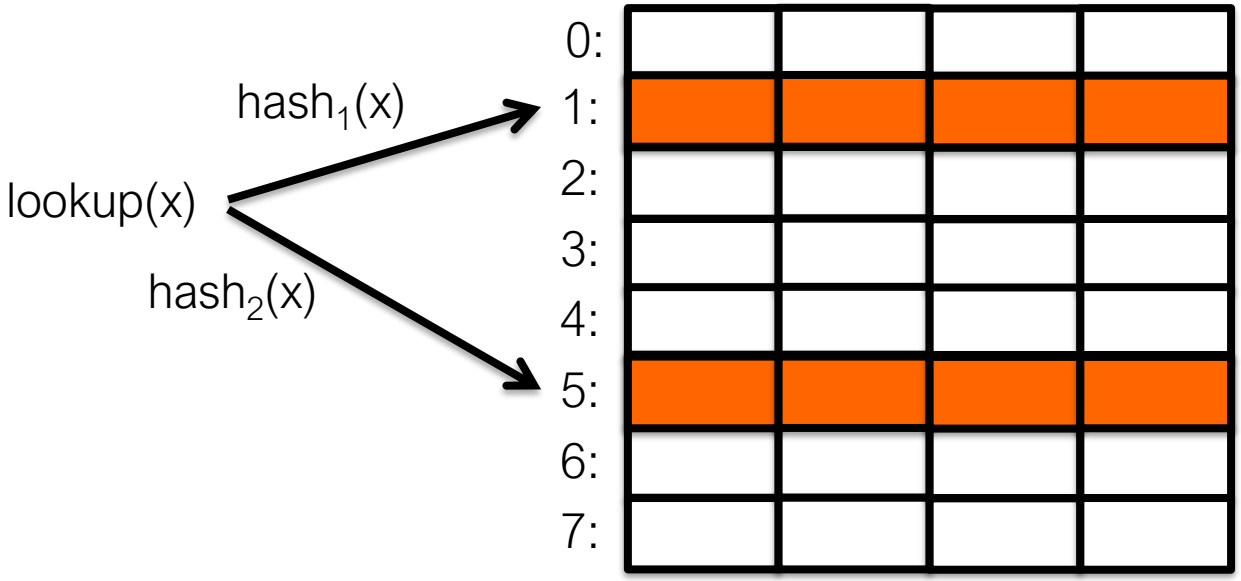

Standard cuckoo hashing doesn't work with fingerprints

[Pagh2004] Cuckoo hashing.

Standard Cuckoo Requires Storing Each Item

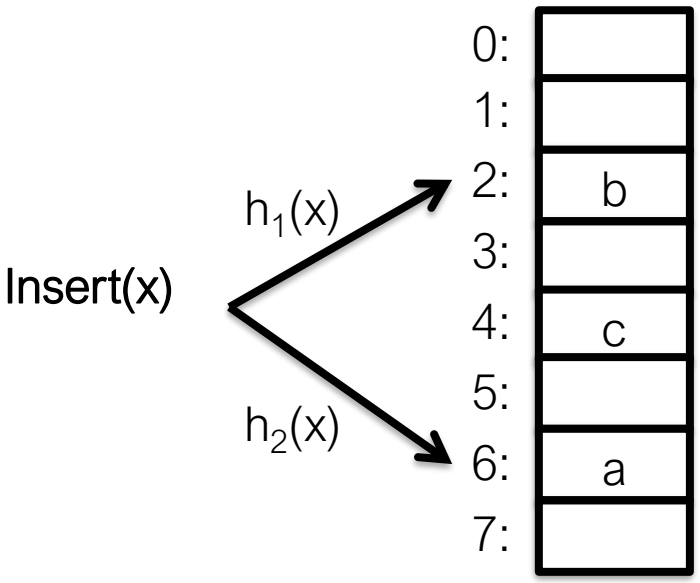

Standard Cuckoo Requires Storing Each Item

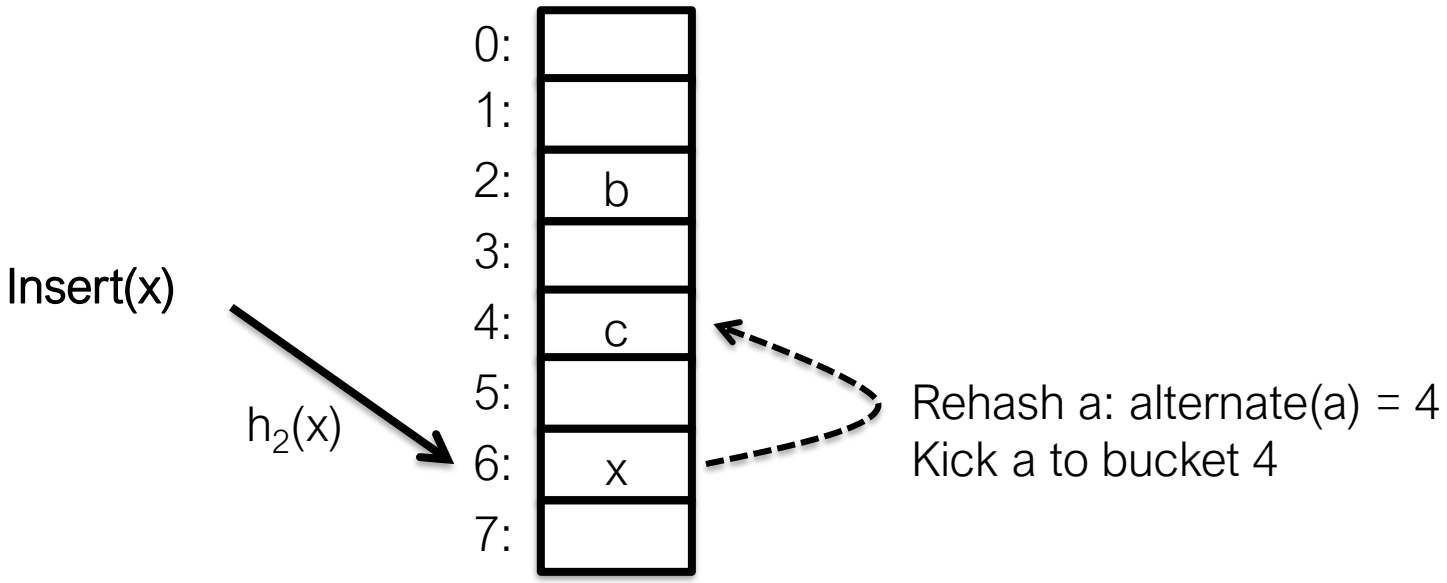

Standard Cuckoo Requires Storing Each Item

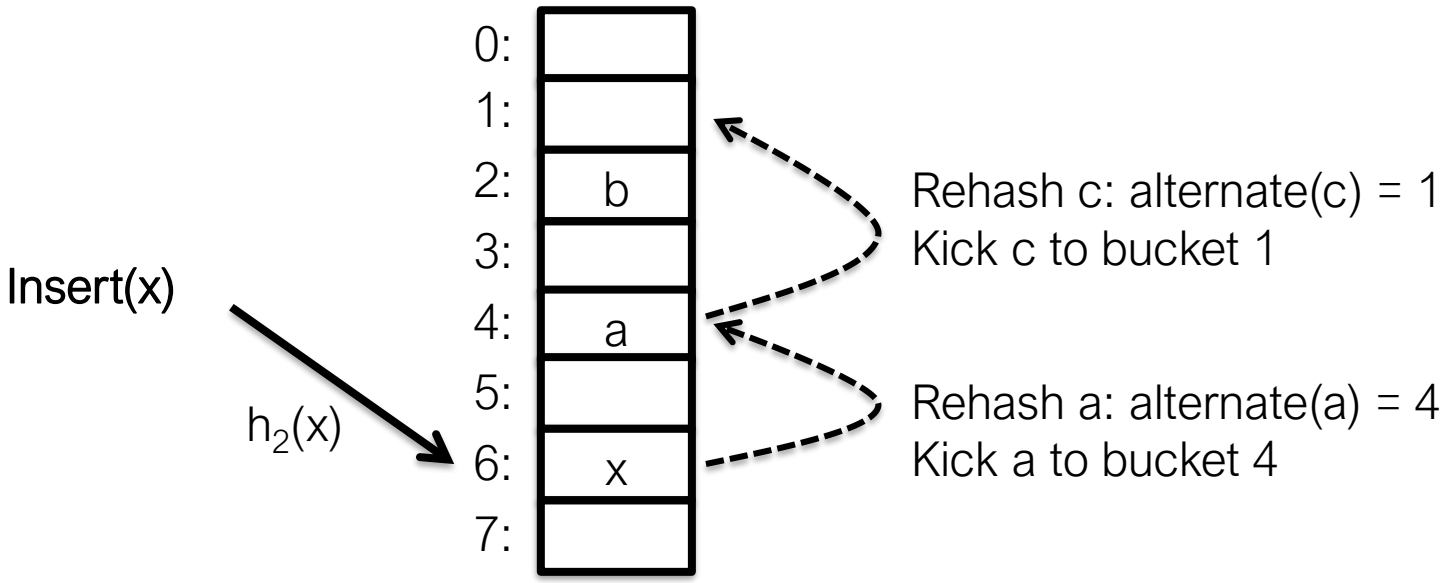

Standard Cuckoo Requires Storing Each Item

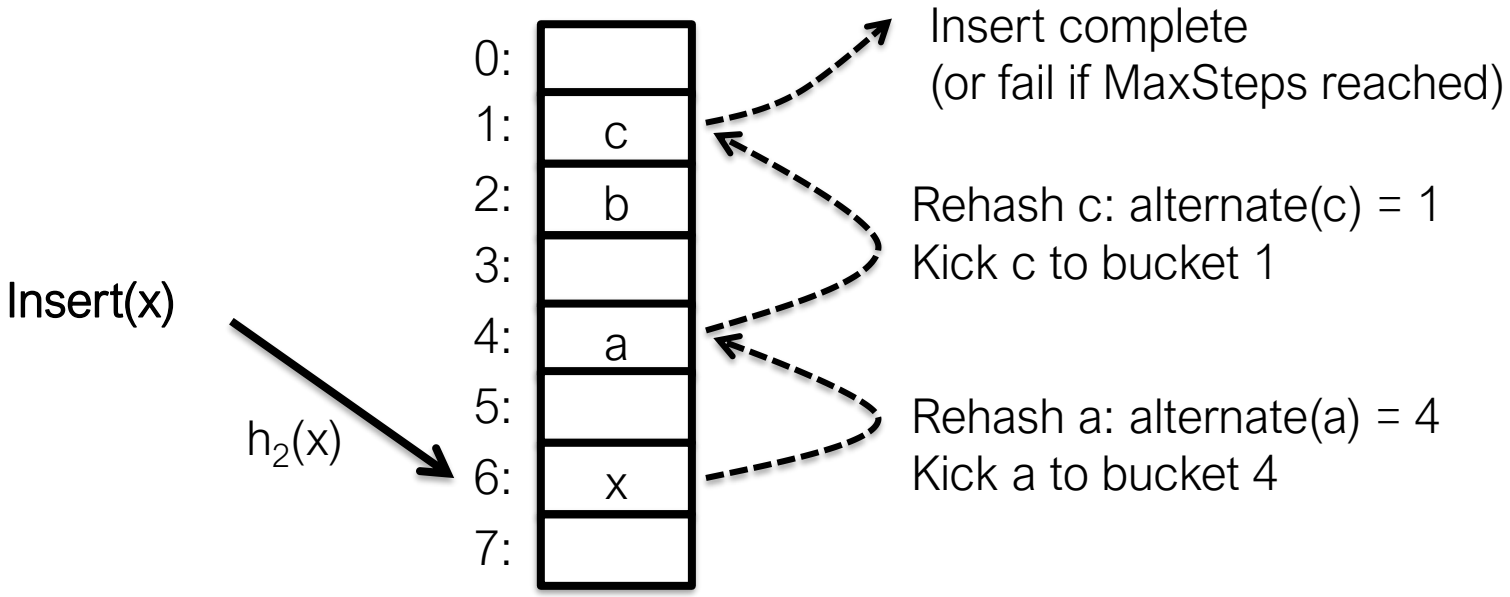

Challenge: How to Perform Cuckoo?

- Cuckoo hashing requires rehashing and displacing existing items

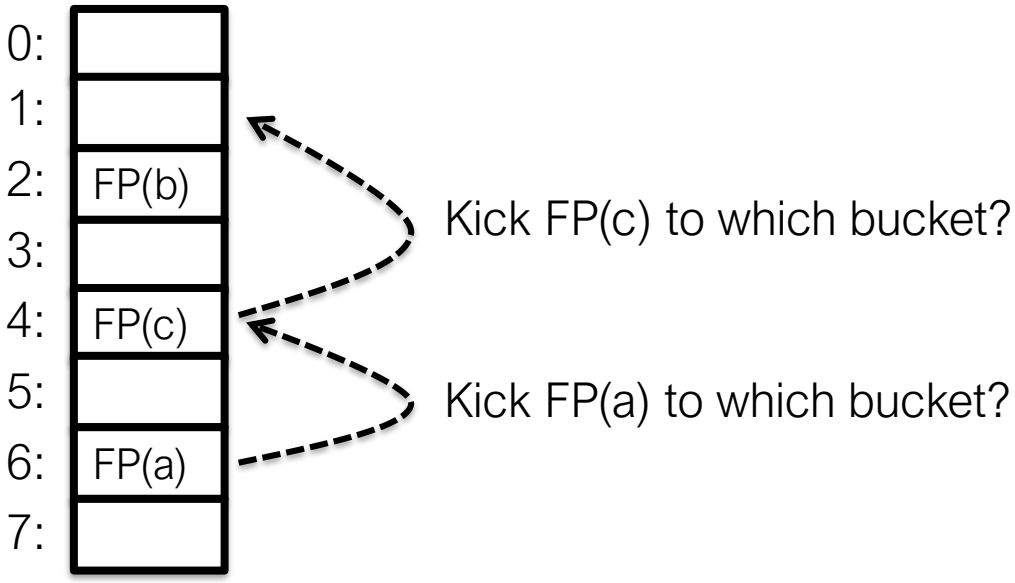

With only fingerprint,
how to calculate item's alternate bucket?

Solution

We Apply Partial-Key Cuckoo

- Standard Cuckoo Hashing: two independent hash functions for two buckets

  $bucket1 = hash_1(x)$

  $bucket2 = hash_2(x)$

- Partial-key Cuckoo Hashing: use one bucket and fingerprint to derive the other [Fan2013]

  $bucket1 = hash(x)$

  $bucket2 = bucket1 \oplus hash(FP(x))$

To displace existing fingerprint:

$alternate(x) = current(x) \oplus hash(FP(x))$

[Fan2013] MemC3: Compact and Concurrent MemCache with Dumber Caching and Smarter Hashing

Solution

Partial Key Cuckoo Hashing

- Perform cuckoo hashing on fingerprints

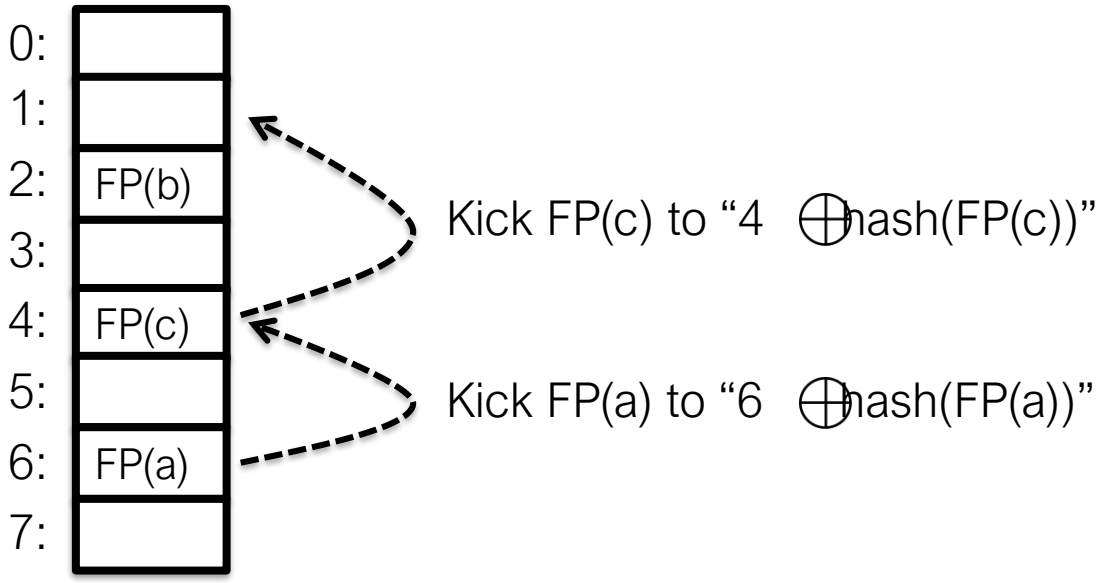

Can we still achieve high space utilization with partial-key cuckoo hashing?

Fingerprints Must Be "Long" for Space Efficiency

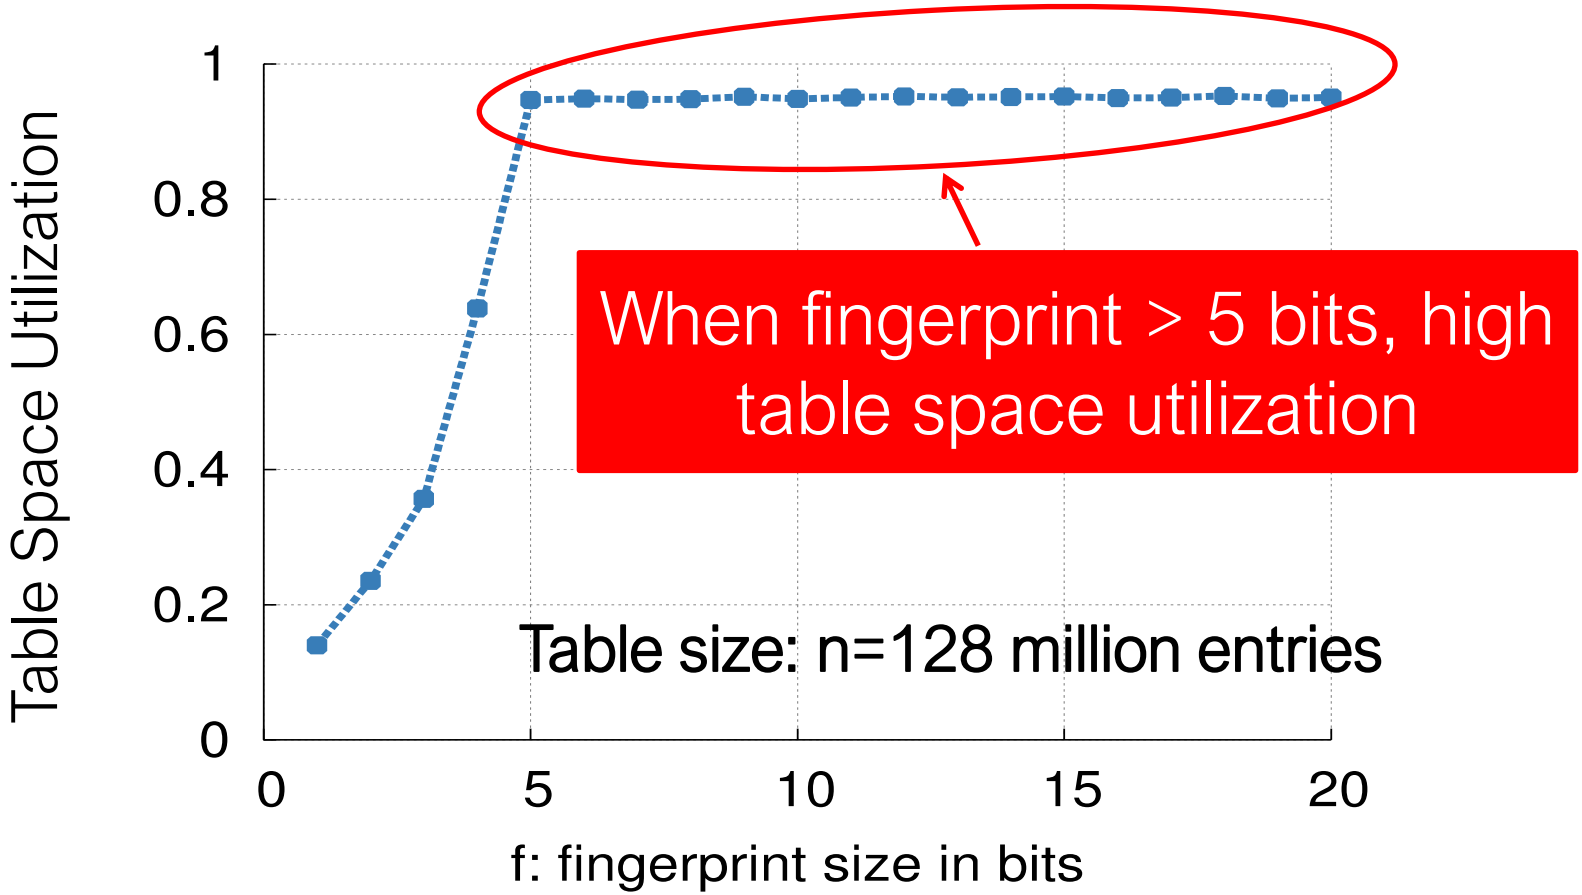

- Fingerprint must be Ω(logn/b) bits in theory
  - n: hash table size, b: bucket size
  - see more analysis in paper

Semi-Sorting: Further Save 1 bit/item

- Based on observation:
  - A monotonic sequence of integers is easier to compress[Bonomi2006]

- Semi-Sorting:
  - Sort fingerprints sorted in each bucket
  - Compress sorted fingerprints

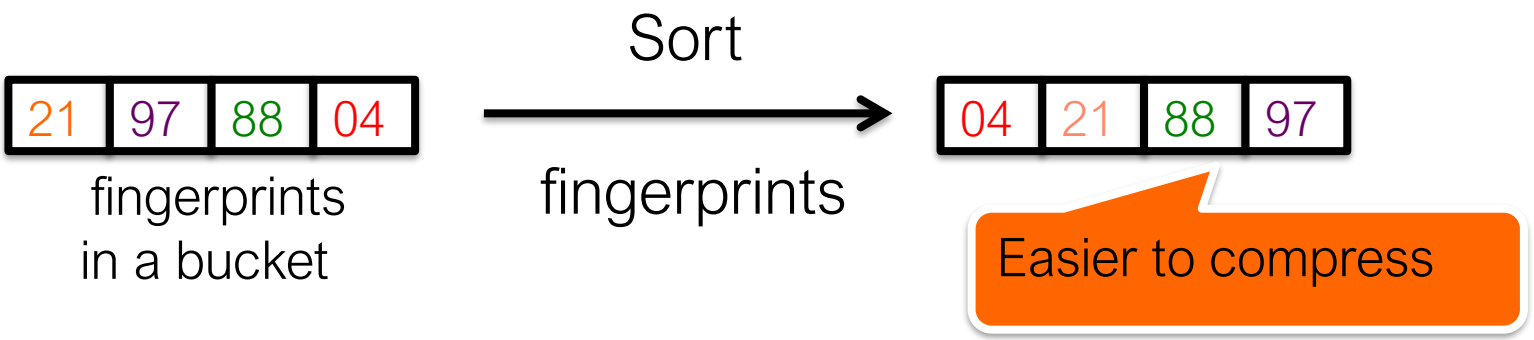

+ For 4-way bucket, save one bit per item

-- Slower lookup / insert

[Bonomi2006] Beyond Bloom filters: From approximate membership checks to ap- proximate state machines.

Space Efficiency

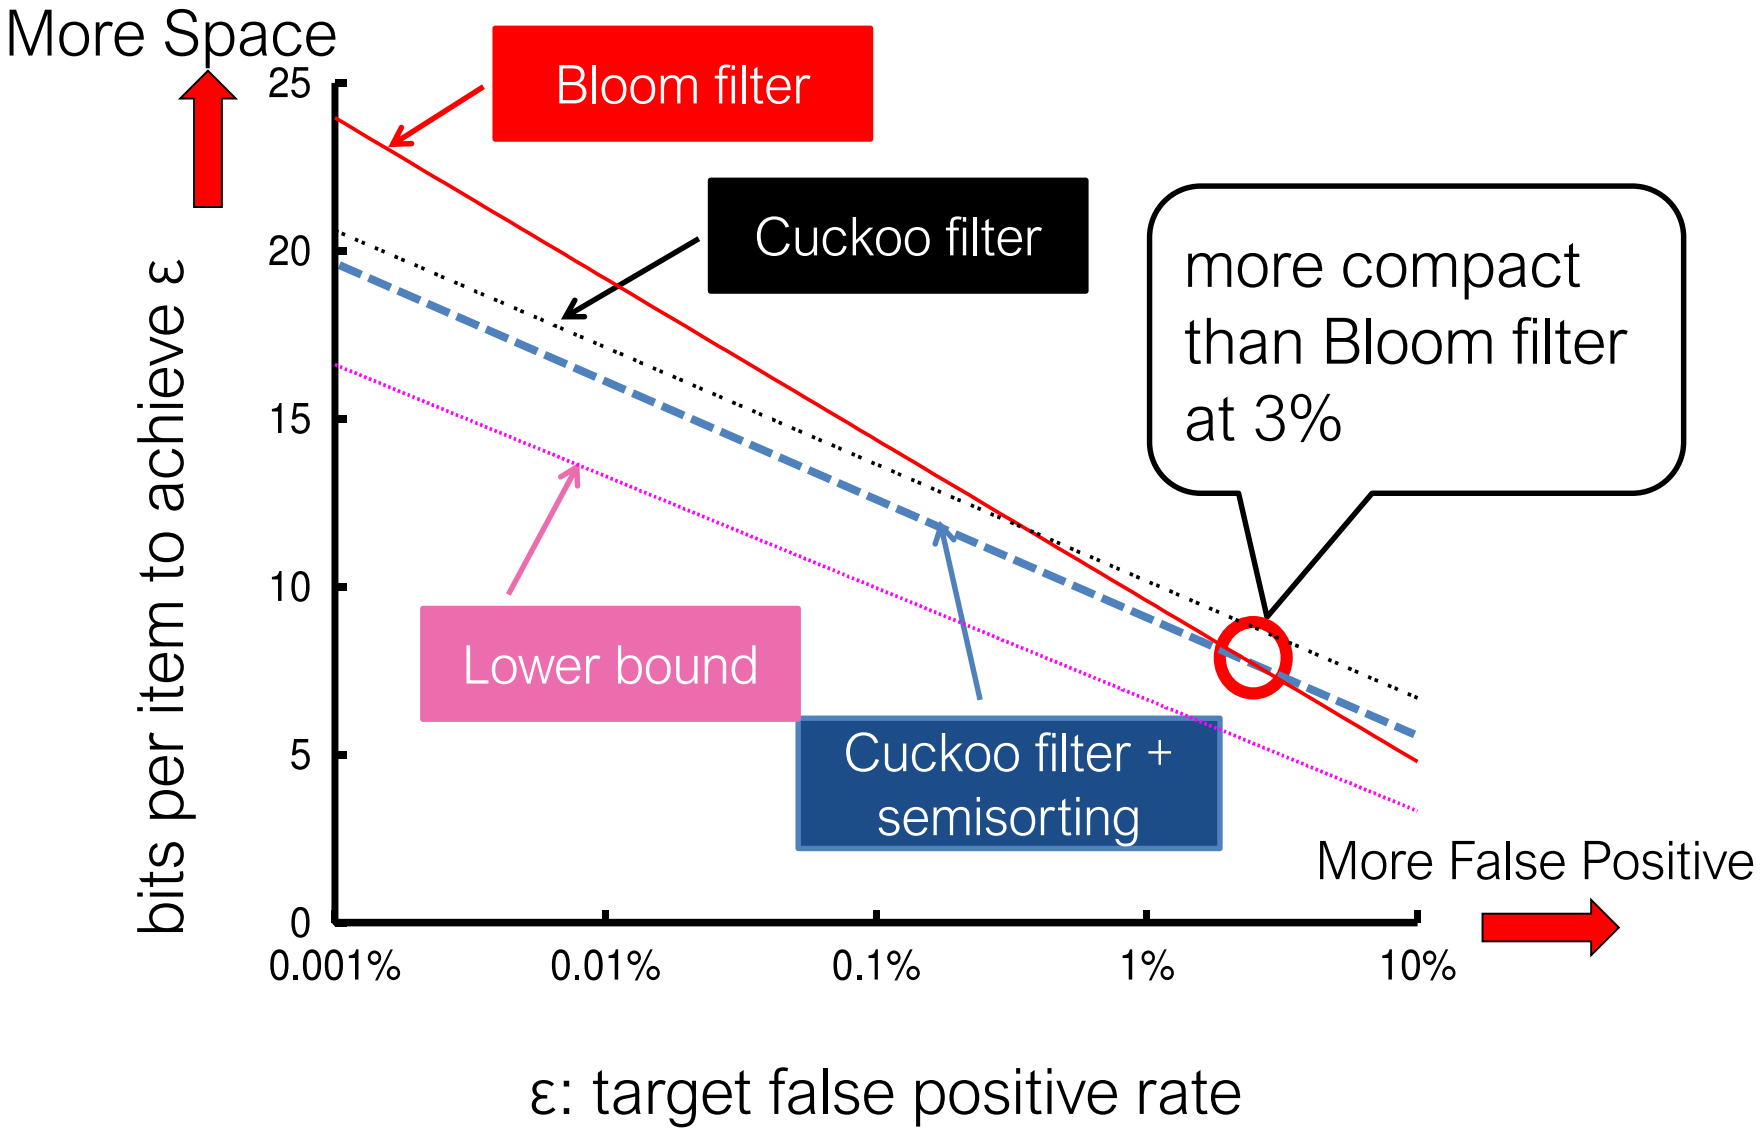

ε: target false positive rate

Outline

- Background
- Cuckoo filter algorithm
- Performance evaluation
- Summary

Evaluation

- Compare cuckoo filter with
  - Bloom filter (cannot delete)
  - Blocked Bloom filter [Putze2007] (cannot delete)
  - d-left counting Bloom filter [Bonomi2006]
  - Cuckoo filter + semisorting
  - More in the paper

- C++ implementation, single threaded

[Putze2007] Cache-, hash- and space- efficient bloom filters.

[Bonomi2006] Beyond Bloom filters: From approximate membership checks to approximate state machines.

Lookup Performance (MOPS)

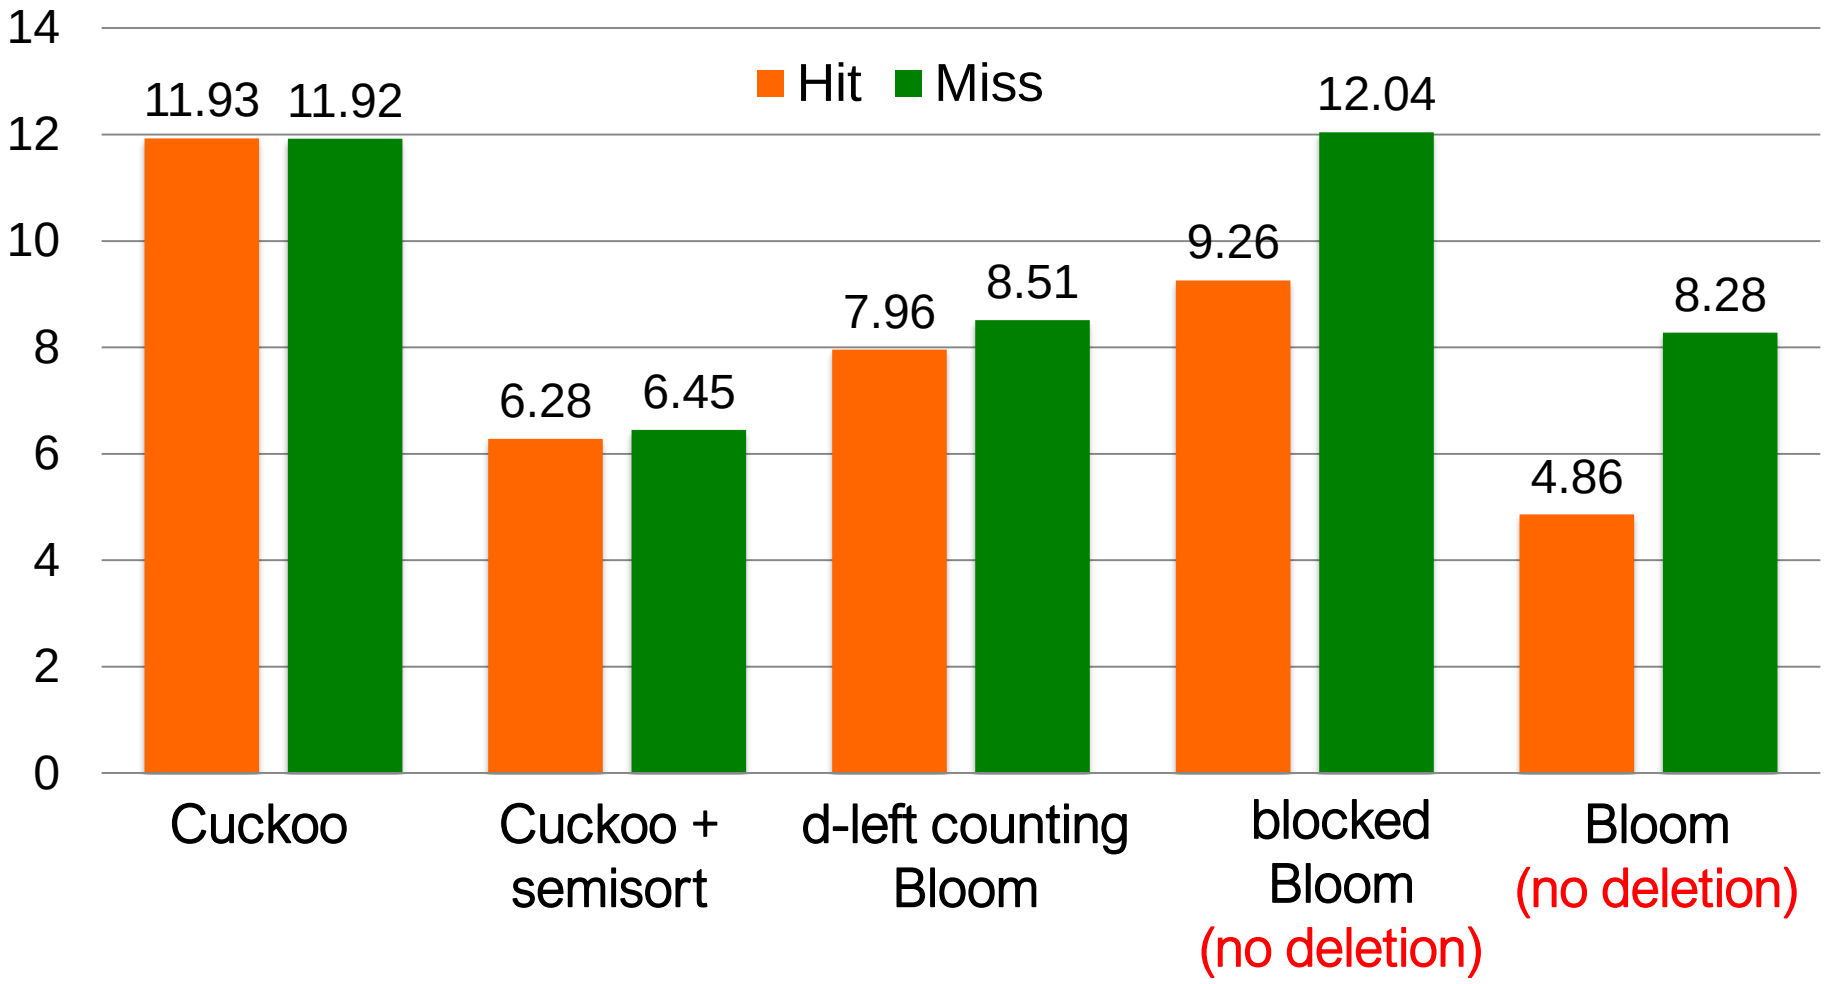

Cuckoo filter is among the fastest regardless workloads.

Insert Performance (MOPS)

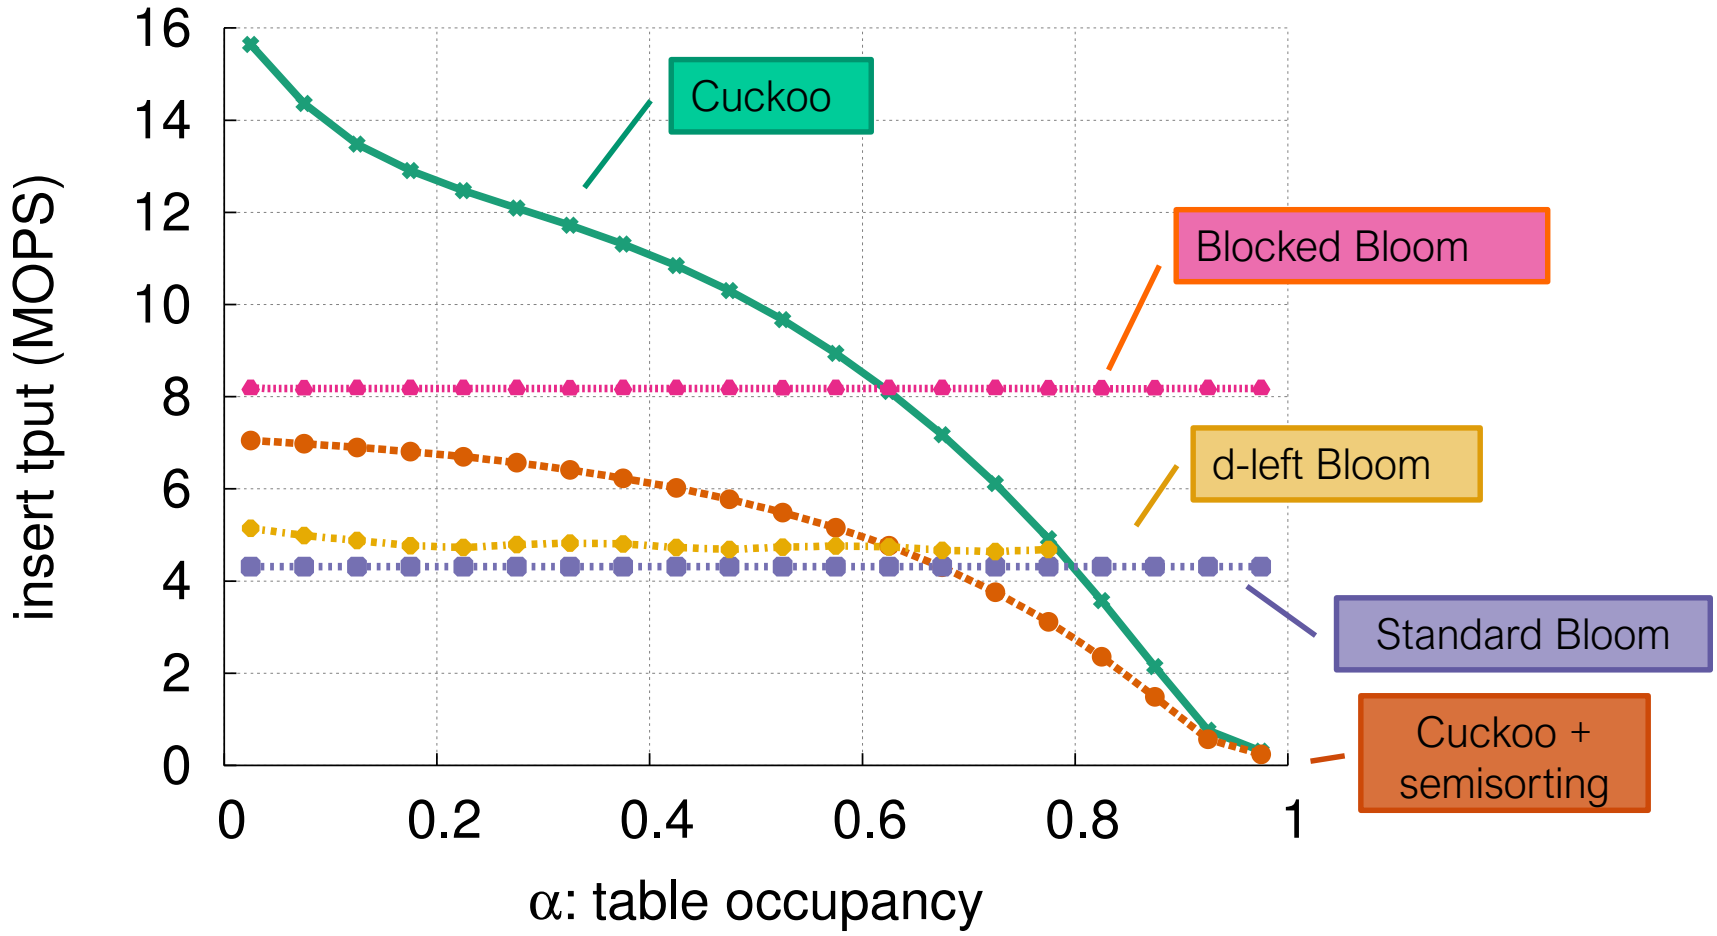

Cuckoo filter has decreasing insert rate, but overall is only slower than blocked Bloom filter.

Summary

- Cuckoo filter, a Bloom filter replacement:
  - Deletion support
  - High performance
  - Less Space than Bloom filters in practice
  - Easy to implement

- Source code available in C++:
  - https://github.com/efficient/cuckoofilter