

 Intro to Database Systems (15-445/645)

# 07 Hash Tables

Carnegie  
Mellon  
University

SPRING  
2023

Charlie  
Garrod

# ADMINISTRIVIA

---

Homework 2 due February 17<sup>th</sup>.

Project 1 due February 19<sup>th</sup>.

→ Saturday office hours: February 18<sup>th</sup> 3-5 p.m.

# COURSE STATUS

We are now going to talk about how to support the DBMS's execution engine to read/write data from pages.

Two types of data structures:

- Hash Tables
- Trees

Query Planning

Operator Execution

Access Methods

Buffer Pool Manager

Disk Manager

# DATA STRUCTURES

---

Internal Meta-data  
Core Data Storage  
Temporary Data Structures  
Table Indexes

# DESIGN DECISIONS

---

## Data Organization

→ How we layout data structure in memory/pages and what information to store to support efficient access.

## Concurrency

→ How to enable multiple threads to access the data structure at the same time without causing problems.

# HASH TABLES

---

A hash table implements an unordered associative array that maps keys to values.

It uses a hash function to compute an offset into this array for a given key, from which the desired value can be found.

Space Complexity:  $O(n)$


Time Complexity:

→ Average:  $O(1)$   ***Databases need to care about constants!***

# STATIC HASH TABLE

Allocate a giant array that has one slot for every element you need to store.

To find an entry, mod the key by the number of elements to find the offset in the array.



# UNREALISTIC ASSUMPTIONS


Assumption#1: Number of elements is known ahead of time and fixed.

Assumption #2: Each key is unique.

Assumption #3: Perfect hash function.

→ If  $\text{key1} \neq \text{key2}$ , then  
 $\text{hash}(\text{key1}) \neq \text{hash}(\text{key2})$

$\text{hash}(\text{key}) \% N$





# HASH TABLE

---

## Design Decision #1: Hash Function

- How to map a large key space into a smaller domain.
- Trade-off between being fast vs. collision rate.

## Design Decision #2: Hashing Scheme

- How to handle key collisions after hashing.
- Trade-off between allocating a large hash table vs. additional instructions to get/put keys.

# TODAY'S AGENDA

---

Hash Functions  
Static Hashing Schemes  
Dynamic Hashing Schemes

# HASH FUNCTIONS

---

For any input key, return an integer representation of that key.

We do not want to use a cryptographic hash function for DBMS hash tables (e.g., SHA-2) .

We want something that is fast and has a low collision rate.

# HASH FUNCTIONS

---

## CRC-64 (1975)

→ Used in networking for error detection.

## MurmurHash (2008)

→ Designed as a fast, general-purpose hash function.

## Google CityHash (2011)

→ Designed to be faster for short keys (<64 bytes).

## Facebook XXHash (2012)

→ From the creator of zstd compression.

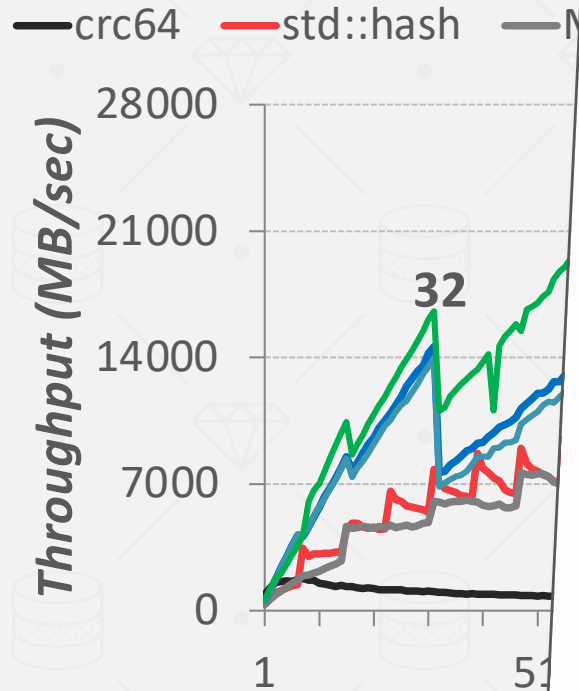
← State-of-the-art

## Google FarmHash (2014)

→ Newer version of CityHash with better collision rates.

# HASH FU

Intel



GitHub repository: [rurban / smhasher](#) Public

Code Issues 26 Pull requests 1 Actions Projects Wiki Security Insights

master [smhasher / README.md](#)

Latest commit 9a5e665 on Aug 19 History

7 contributors

444 lines (395 sloc) 37.4 KB

## SMhasher

build passing build failing build failing

| Hash function       | MiB/sec     | cycl./hash | cycl./map  | size | Quality problems                         |
|---------------------|-------------|------------|------------|------|--|
| donothing32         | 15316474.36 | 6.00       | -          | 13   | bad seed 0, test NOP                     |
| donothing64         | 15330019.19 | 6.00       | -          | 13   | bad seed 0, test NOP                     |
| donothing128        | 15278983.09 | 6.00       | -          | 13   | bad seed 0, test NOP                     |
| NOP_OAAT_read64     | 28467.50    | 18.48      | -          | 47   | test NOP                                 |
| BadHash             | 524.81      | 96.20      | -          | 47   | bad seed 0, test FAIL                    |
| sumhash             | 7169.08     | 27.12      | -          | 363  | bad seed 0, test FAIL                    |
| sumhash32           | 22556.18    | 22.98      | -          | 863  | UB, test FAIL                            |
| multiply_shift      | 5418.36     | 28.69      | 157.11 (3) | 345  | bad seeds & 0xffffffff, fails most tests |
| pair_multiply_shift | 3716.95     | 40.22      | 186.34 (3) | 609  | fails most tests                         |

Key Size (bytes)

Source: [Fredrik Widlund](#)

# STATIC HASHING SCHEMES

---

**Approach #1: Linear Probe Hashing**

**Approach #2: Robin Hood Hashing**

**Approach #3: Cuckoo Hashing**

# LINEAR PROBE HASHING

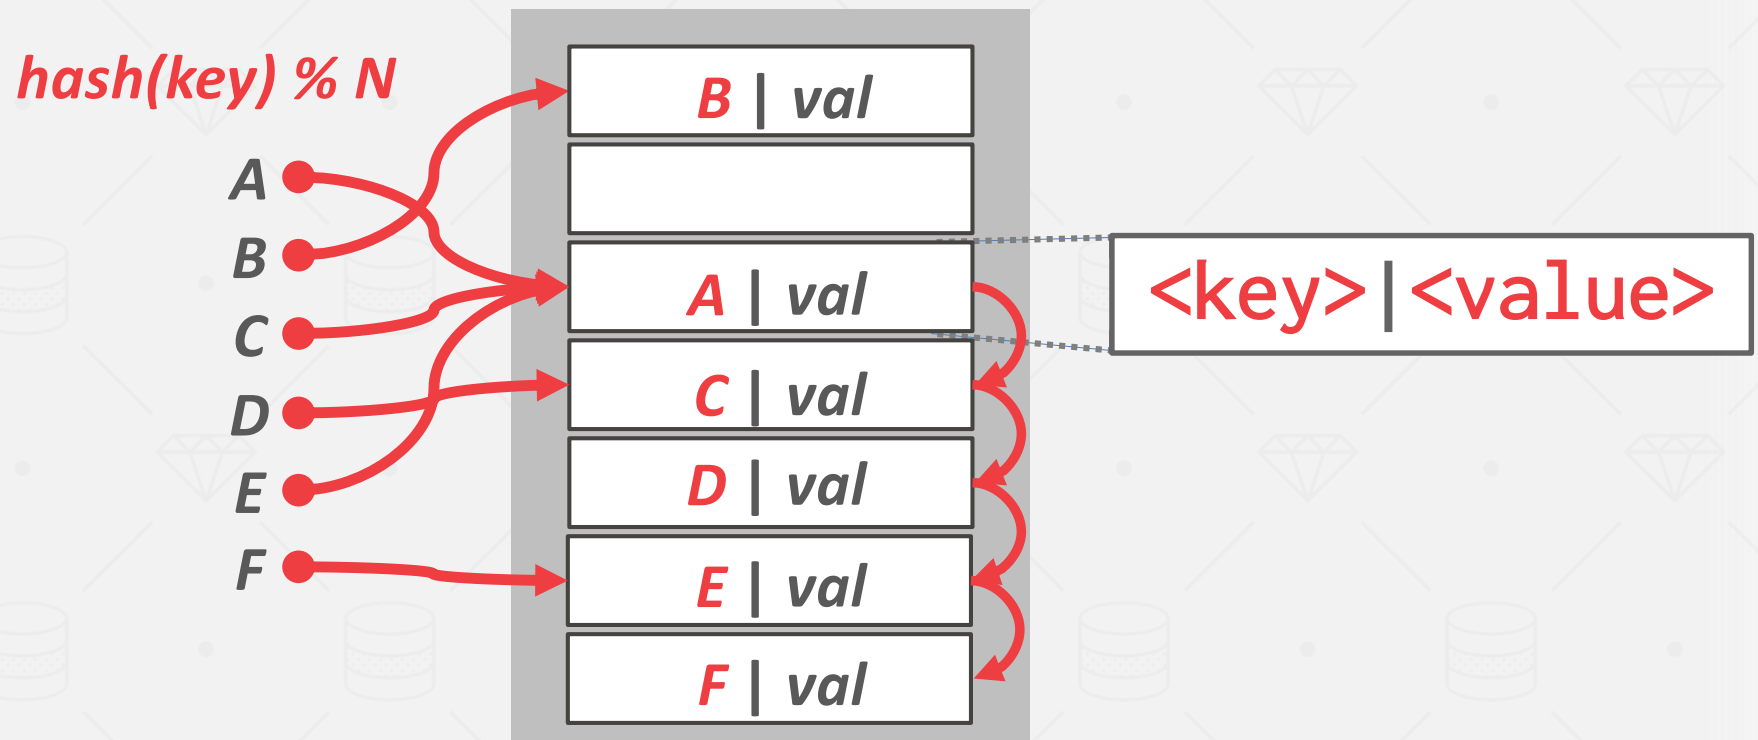
---

Single giant table of slots.

Resolve collisions by linearly searching for the next free slot in the table.

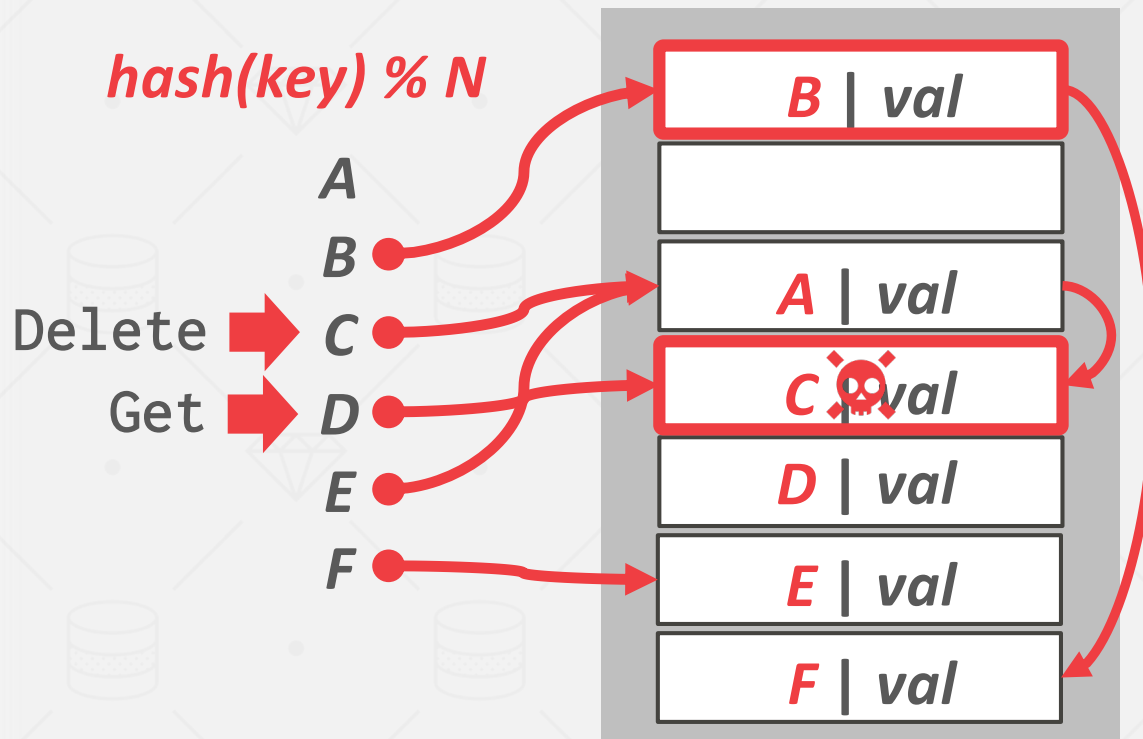
- To determine whether an element is present, hash to a location in the index and scan for it.
- Must store the key in the index to know when to stop scanning.
- Insertions and deletions are generalizations of lookups.

# LINEAR PROBE HASHING





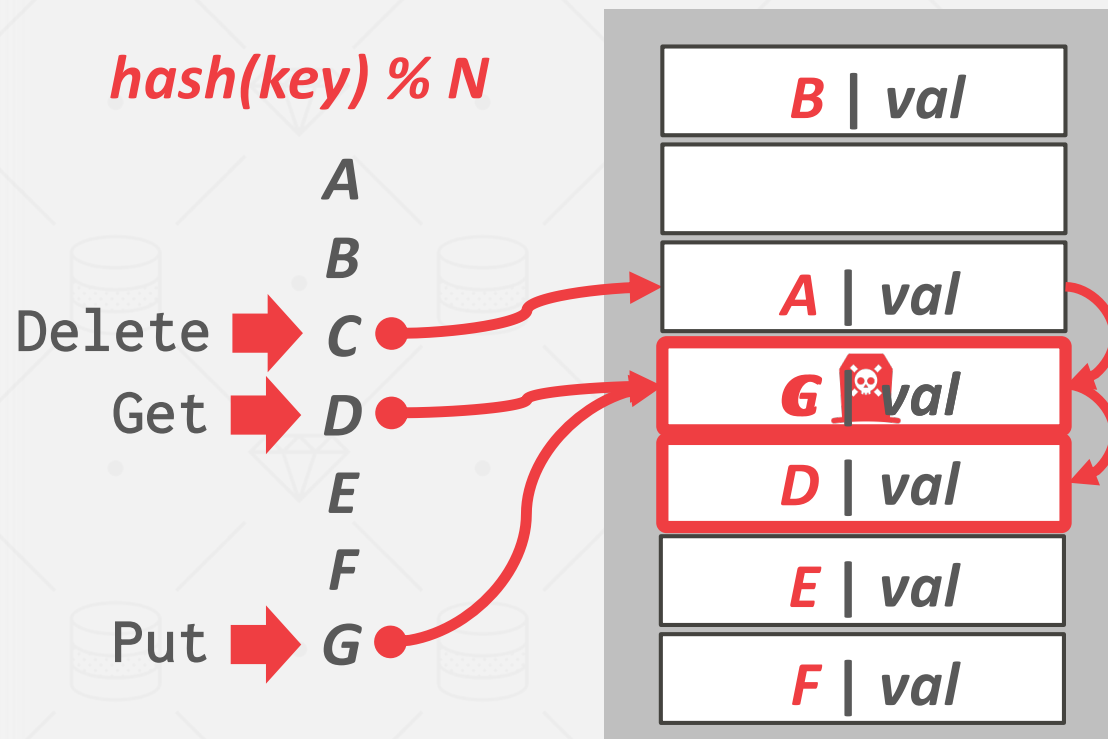
# LINEAR PROBE HASHING – DELETES



## Approach #1: Movement

- Rehash keys until you find the first empty slot.
- Nobody actually does this.

# LINEAR PROBE HASHING – DELETES




## Approach #2: Tombstone

- Set a marker to indicate that the entry in the slot is logically deleted.
- You can reuse the slot for new keys.
- May need periodic garbage collection.

# NON-UNIQUE KEYS

## Choice #1: Separate Linked List


→ Store values in separate storage area for each key.



## Choice #2: Redundant Keys

→ Store duplicate keys entries together in the hash table.

→ This is what most systems do.



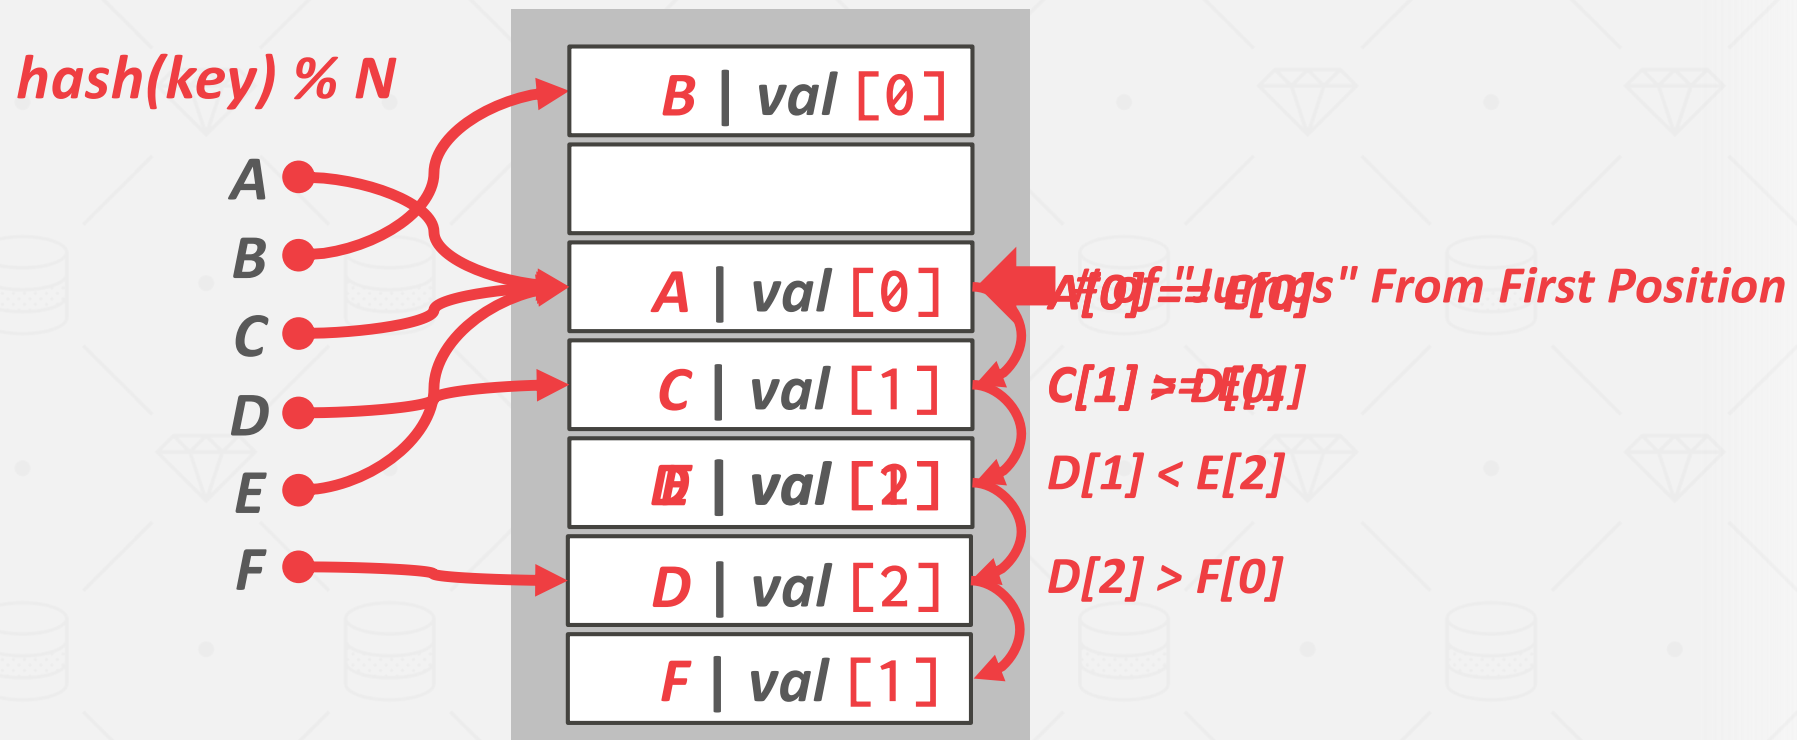
# ROBIN HOOD HASHING

---

Variant of linear probe hashing that steals slots from "rich" keys and give them to "poor" keys.

- Each key tracks the number of positions they are from where its optimal position in the table.
- On insert, a key takes the slot of another key if the first key is farther away from its optimal position than the second key.

# ROBIN HOOD HASHING



# CUCKOO HASHING

---

Use multiple hash tables with different hash function seeds.


- On insert, check every table and pick anyone that has a free slot.
- If no table has a free slot, evict the element from one of them and then re-hash it find a new location.

Look-ups and deletions are always  **$O(1)$**  because only one location per hash table is checked.

Best [open-source implementation](#) is from CMU.

# CUCKOO HASHING

Hash Table #1



Put A

$hash_1(A)$   $hash_2(A)$

Put B

$hash_1(B)$   $hash_2(B)$

Put C

$hash_1(C)$   $hash_2(C)$


$hash_1(B)$

$hash_2(A)$

Get B

$hash_1(B)$   $hash_2(B)$

Hash Table #2



## LIES, LIES, AND ...

---

I think we're lying: cuckoo hashing does not actually use multiple hash tables.

→ The essence is that multiple hash functions map a key to different slots (in a single hash table).

`hash(key) % N` is not a good way to map a key within the size (N) of the hash table.

→ Make N a power of 2 and bit-mask instead.

Hashing is not actually  $O(1)$ ...



## OBSERVATION

---

The previous hash tables require the DBMS to know the number of elements it wants to store.

→ Otherwise, it must rebuild the table if it needs to grow/shrink in size.

Dynamic hash tables incrementally resize themselves as needed.

- Chained Hashing
- Extendible Hashing
- Linear Hashing

# CHAINED HASHING


---

Maintain a linked list of buckets for each slot in the hash table.

Resolve collisions by placing all elements with the same hash key into the same bucket.

- To determine whether an element is present, hash to its bucket and scan for it.
- Insertions and deletions are generalizations of lookups.

# CHAINED HASHING



# EXTENDIBLE HASHING

---

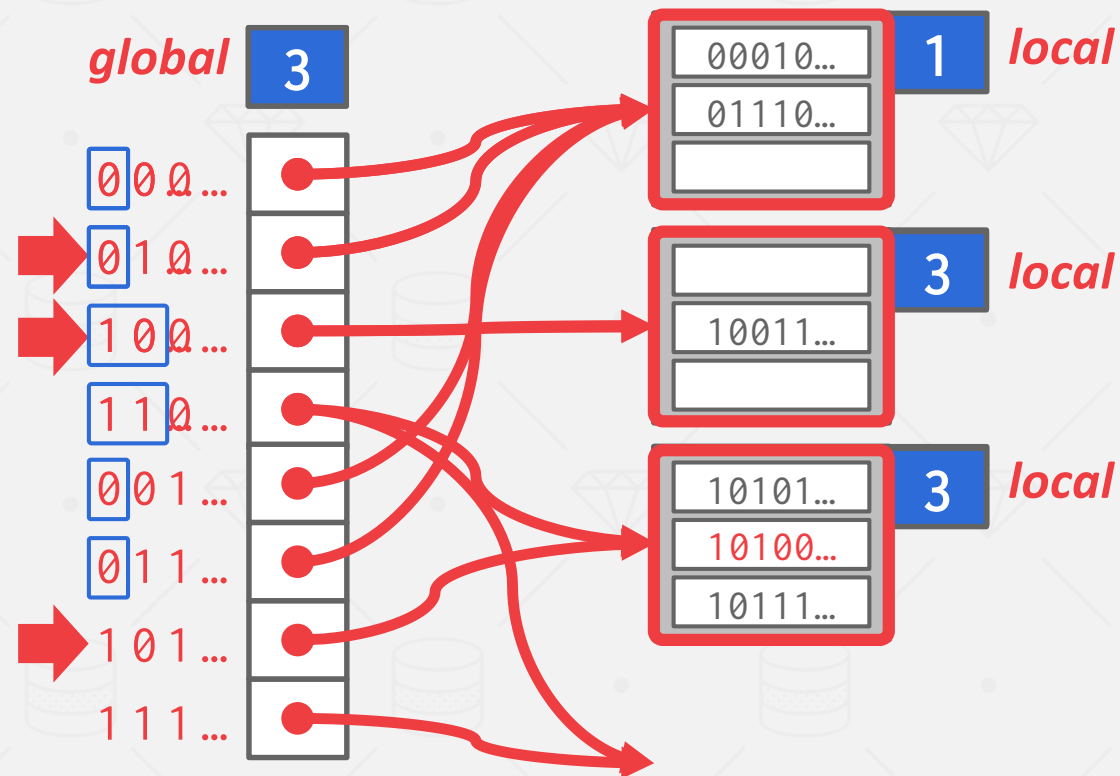
Chained-hashing approach where we split buckets instead of letting the linked list grow forever.

Multiple slot locations can point to the same bucket chain.

Reshuffle bucket entries on split and increase the number of bits to examine.

→ Data movement is localized to just the split chain.

# EXTENDIBLE HASHING



Get A  
 $hash(A) = 01110...$

Put B  
 $hash(B) = 10111...$

Put C  
 $hash(C) = 10100...$

# LINEAR HASHING

---

The hash table maintains a pointer that tracks the next bucket to split.

→ When any bucket overflows, split the bucket at the pointer location.

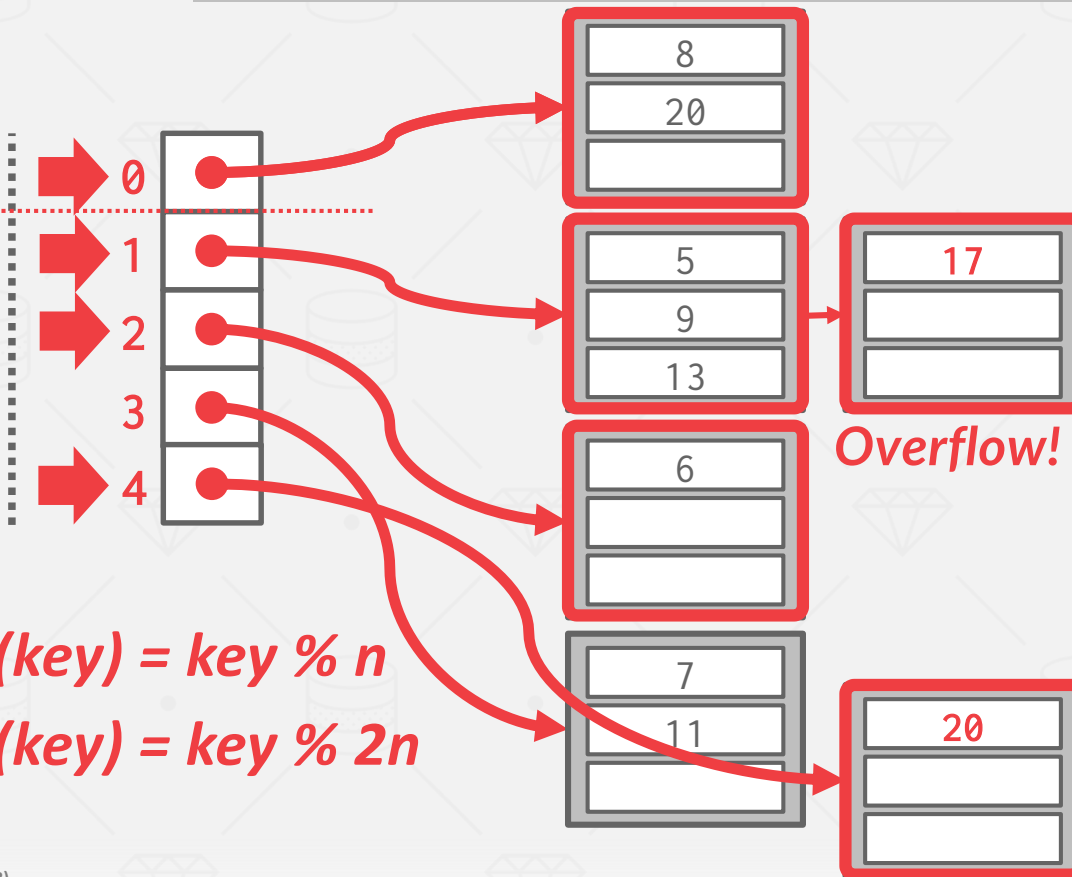
Use multiple hashes to find the right bucket for a given key.

Can use different overflow criterion:

- Space Utilization
- Average Length of Overflow Chains

# LINEAR HASHING

*Split  
Pointer*



Get 6

$$\text{hash}_1(6) = 6 \% 4 = 2$$

Put 17

$$\text{hash}_1(17) = 17 \% 4 = 1$$

Get 20

$$\text{hash}_1(20) = 20 \% 4 = 0$$

$$\text{hash}_2(20) = 20 \% 8 = 4$$

Get 9

$$\text{hash}_1(9) = 9 \% 4 = 1$$

$$\text{hash}_1(\text{key}) = \text{key} \% n$$

$$\text{hash}_2(\text{key}) = \text{key} \% 2n$$

# LINEAR HASHING

---

Splitting buckets based on the split pointer will eventually get to all overflowed buckets.

→ When the pointer reaches the last slot, delete the first hash function and move back to beginning.



# CONCLUSION

---

Fast data structures that support  **$O(1)$**  look-ups that are used all throughout DBMS internals.  
→ Trade-off between speed and flexibility.

Hash tables are usually **not** what you want to use for a table index...

# NEXT CLASS

---

## B+ Trees

→ aka "The Greatest Data Structure of All Time"