

Administrivia

- Assignment 3 released
- Ten students will present their project in class.

Hash Tables

Hash Tables

- A hash table implements an unordered associative array that maps keys to values.
 - ▶ `mymap.insert('a', 50);`
 - ▶ `mymap['b']=100;`
 - ▶ `mymap.find('a')`
 - ▶ `mymap['a']`
- It uses a hash function to compute an offset into the array for a given key, from which the desired value can be found.


Assumptions

- You know the number of elements ahead of time.
- Each key is unique (*e.g.*, SSN ID \longrightarrow Name).
- Perfect hash function (no **collision**).
 - If $\text{key1} \neq \text{key2}$, then $\text{hash}(\text{key1}) \neq \text{hash}(\text{key2})$


Hash Table: Design Decisions

- Design Decision 1: **Hash Function**
 - How to map a large key space into a smaller domain of array offsets.
 - 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 steps to find/insert keys.

Linear Probe Hashing




Linear Probe Hashing



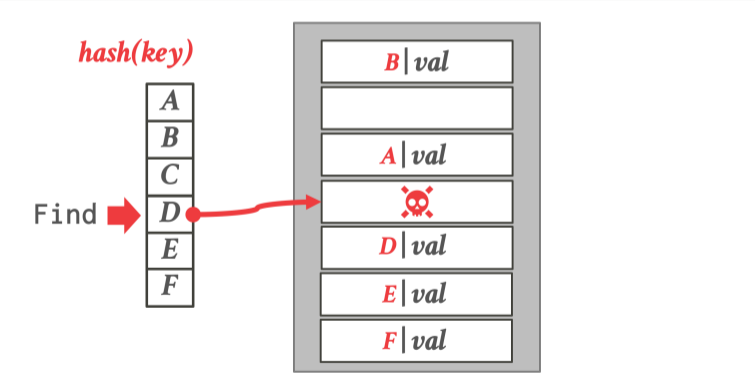
Linear Probe Hashing – Delete

- It is **not sufficient** to simply delete the key.
- This would affect searches for other keys that have a hash value earlier than the emptied cell, but that are stored in a position later than the emptied cell.
- Solutions:
 - ▶ Approach 1: Tombstone
 - ▶ Approach 2: Movement


Linear Probe Hashing – Delete




Linear Probe Hashing – Delete



Linear Probe Hashing – Delete




Linear Probe Hashing – Delete



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.



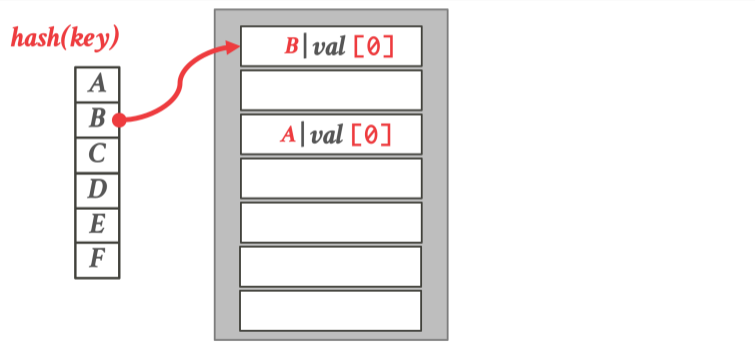
Robin Hood Hashing

- 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.

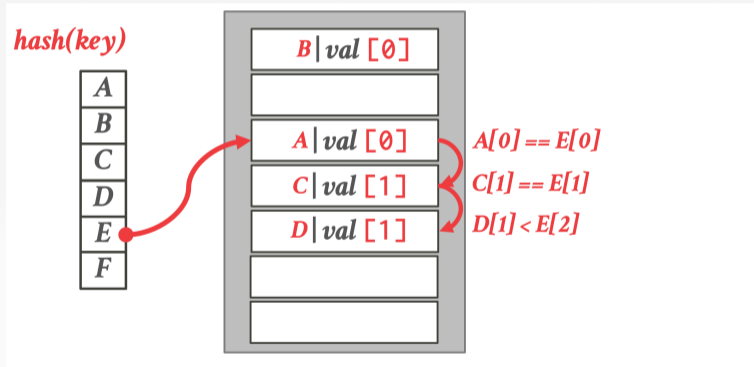
Robin Hood Hashing




Robin Hood Hashing




Robin Hood Hashing




Robin Hood Hashing



Robin Hood Hashing




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.


Cuckoo Hashing




Cuckoo Hashing



Cuckoo Hashing



Cuckoo Hashing



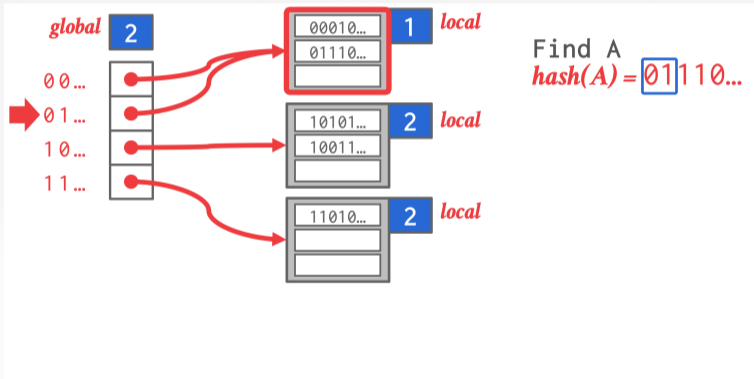
Observation

- Static hashing schemes require the DBMS to know the number of keys to be stored.
 - ▶ Otherwise it has to rebuild the table if it needs to grow/shrink the table in size. Why?
 - ▶ You would have to take a latch on the entire hash table to prevent threads from adding new entries.
- Dynamic hashing schemes resize themselves on demand.
 - ▶ Approach 1: Chained Hashing
 - ▶ Approach 2: Extendible Hashing
 - ▶ Approach 3: Linear Hashing


Chained Hashing

- Unlike static hashing schemes, two different keys may hash to the same offset
- If you want to enforce **unique keys**, then you have perform an additional comparison of each key to determine whether they exactly match
- So, unlike static hashing schemes, need to retain the original key in the table

Extendible Hashing



Linear Hashing



Linear Hashing

