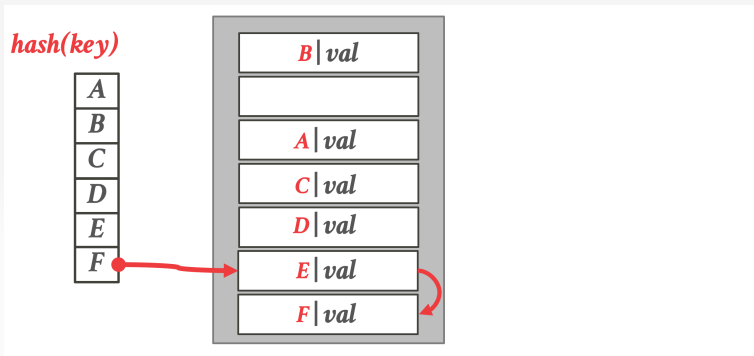
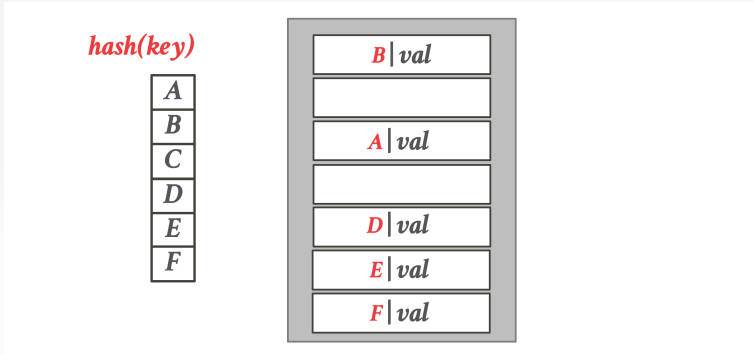


Hash Functions


Linear Probe Hashing



Linear Probe Hashing – Delete




Linear Probe Hashing – Delete




Non-Unique Keys

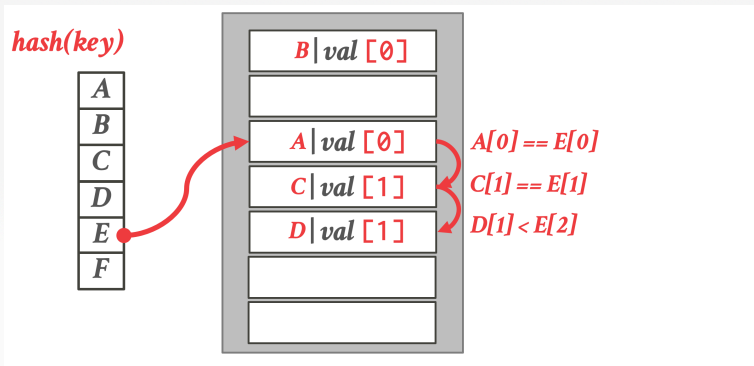
- Choice 1: Separate Linked List
 - Store values in separate storage area for each key.
- Choice 2: Redundant Keys
 - Store duplicate keys entries together in the hash table.




Robin Hood Hashing



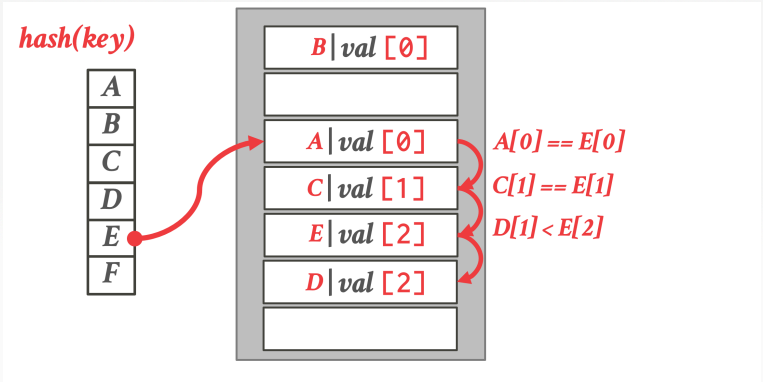
Robin Hood Hashing




Robin Hood Hashing




Robin Hood Hashing



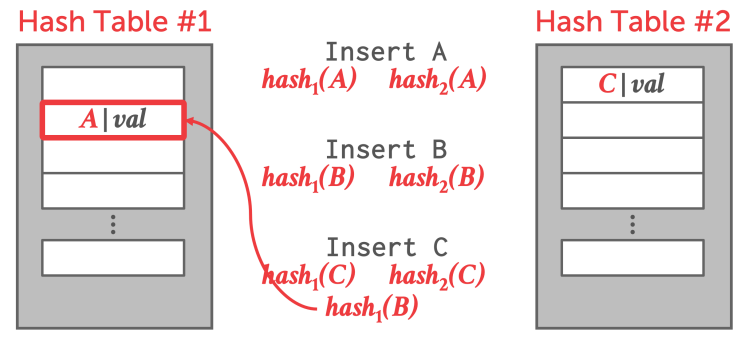
Robin Hood Hashing



Cuckoo Hashing




Cuckoo Hashing



Cuckoo Hashing

Hash Table #1




Insert A
 $hash_1(A) \quad hash_2(A)$

Insert B
 $hash_1(B) \quad hash_2(B)$

Insert C
 $hash_1(C) \quad hash_2(C)$
 $hash_1(B)$
 $hash_2(A)$

Hash Table #2




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

- Maintain a linked list of **buckets for each slot** in the hash table.
- Resolve collisions by placing all keys with the same hash value 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



Chained Hashing

- The hash table can grow infinitely because you just keep adding new buckets to the linked list.
- You only need to take a latch on the bucket to store a new entry or extend the linked list.


Extendible Hashing

- The slot array maps hashes to buckets.
- A hash value may occupy an arbitrary number of bits.
- With extendible hashing, the number of bits that the hash table uses to map hashes to buckets changes over time.
 - ▶ Global counter keeps track of the number of bits that the the hash table uses.
 - ▶ Local counter in each bucket tracks the number of hash bits used by that bucket.


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



Linear Hashing



Linear Hashing vs. Extendible Hashing

- Moving from $hash_i$ to $hash_{i+1}$ in Linear Hashing corresponds to
- Bumping up the global counter in Extendible Hashing
- Linear Hashing
 - ▶ Directory is gradually doubled over the course of a round
 - ▶ A directory can be avoided by a clever choice of the buckets to split
 - ▶ More flexibility: need not always split the appropriate dense bucket

