

# FAST LOOKUP WITH HASHTABLES

---


Problem Solving with Computers-II

STD::UNDORDERED\_SET  
STD::UNORDERED\_MAP

The image shows the C++ logo in blue, followed by a snippet of C++ code. The code is written in a monospaced font and includes a header, namespace declaration, and a main function that prints a message and returns 0.

```
include <iostream>  
using namespace std;  
  
int main(){  
    cout<<"Hola Facebook!n";  
    return 0;  
}
```


# Hash table: Practical Applications



## Network security

- Routers process 2-3 million packets per second
- Security strategies include:
  - Dropping packets from blacklisted sources
  - Need to lookup IP address in blacklist

# Hash table: Practical Applications



Amazon makes approx. \$1.3 billion per day.

In July 2020, Amazon sites had 213 million unique visitors in the US.

## Network security

- Routers process 2-3 million packets per second
- Need efficient lookup of IP address in a blacklist

## Website analytics

- How many unique visitors to a website?
- Need efficient way to de-duplicate entries in very large logs

Many applications rely on fast lookup in an evolving dataset

## Last time: Balanced BST

```
set<string> groceries;  
groceries.insert("Banana") ;  
groceries.insert("Apple");  
groceries.insert("Milk");  
groceries.insert("Bread");
```

**std::set**

**stores unique keys**

```
map<string, int> groceries;  
groceries["Banana"] = 2;  
groceries["Apple"] = 1;  
groceries["Milk"] = 3;  
groceries["Bread"] = 5;
```

**std::map**

**stores key, value pairs**

# Today: Hash table

## Hash table

```
unordered_set<string> groceries;  
groceries.insert("Banana");  
groceries.insert("Apple");  
groceries.insert("Milk");  
groceries.insert("Bread");
```

**std::unordered\_set**  
stores unique keys

```
unordered_map<string, int> groceries;  
groceries["Banana"] = 2;  
groceries["Apple"] = 1;  
groceries["Milk"] = 3;  
groceries["Bread"] = 5;
```

**std::unordered\_map**  
stores key, value pairs

**Operations: find, insert, erase: all  $O(1)$**

- \* not a worst case guarantee
- \* only if the hash table is implemented properly

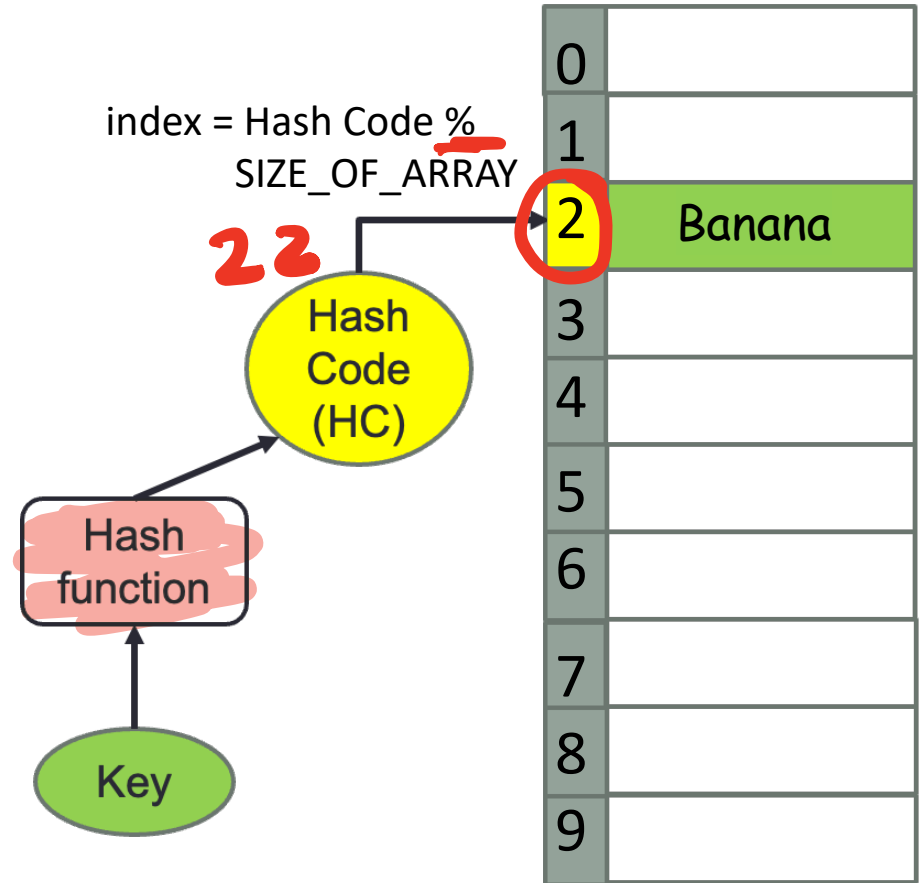
# Hash table: an array with positions indexed by keys

```
unordered_set<int> groceries;  
groceries.insert(2);  
groceries.insert(1);  
groceries.insert(7);  
groceries.insert(3);  
groceries.find(2);
```

0	
1	1
2	Key 2
3	3
4	
5	
6	
7	7
8	
9	

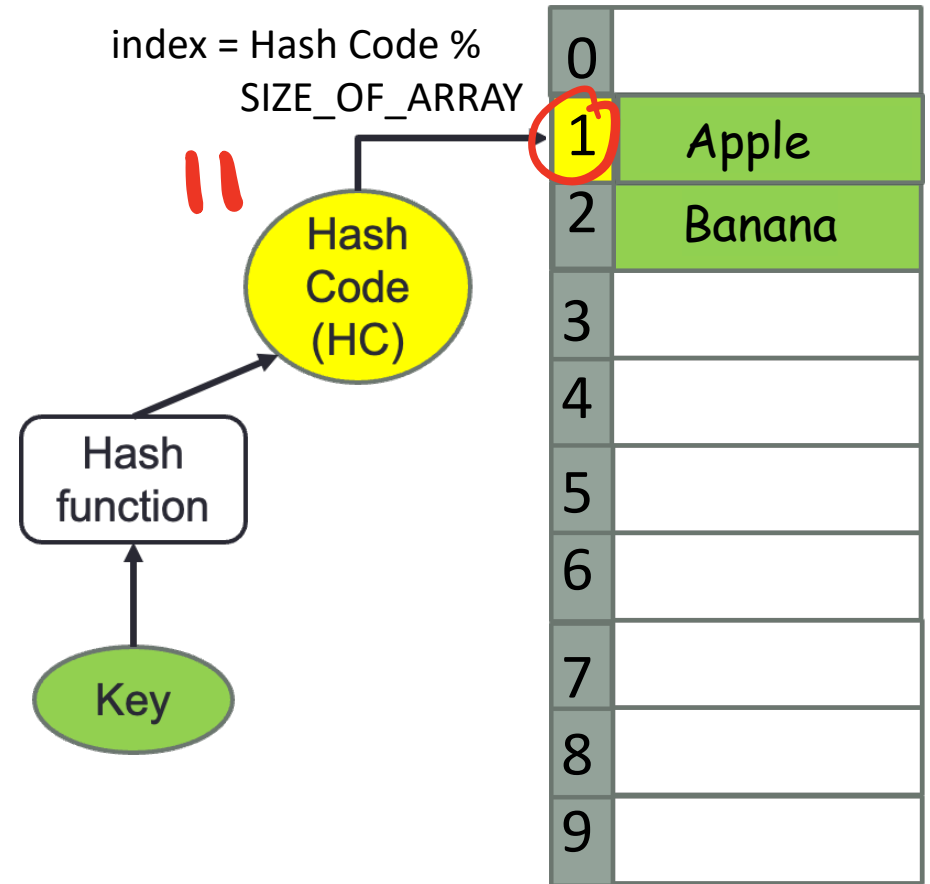
# Hash table: an array with positions indexed by keys

```
unordered_set<string> groceries;  
groceries.insert("Banana");
```



# Hash table is just an array with positions indexed by keys

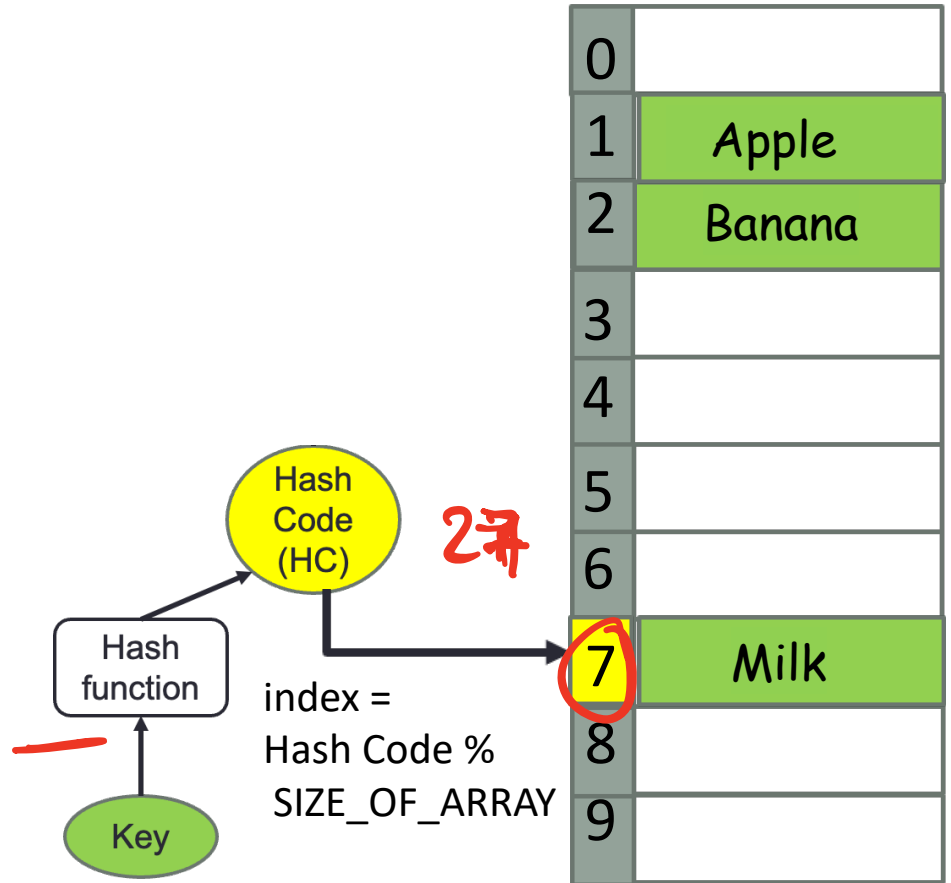
```
unordered_set<string> groceries;  
groceries.insert("Banana");  
groceries.insert("Apple");
```





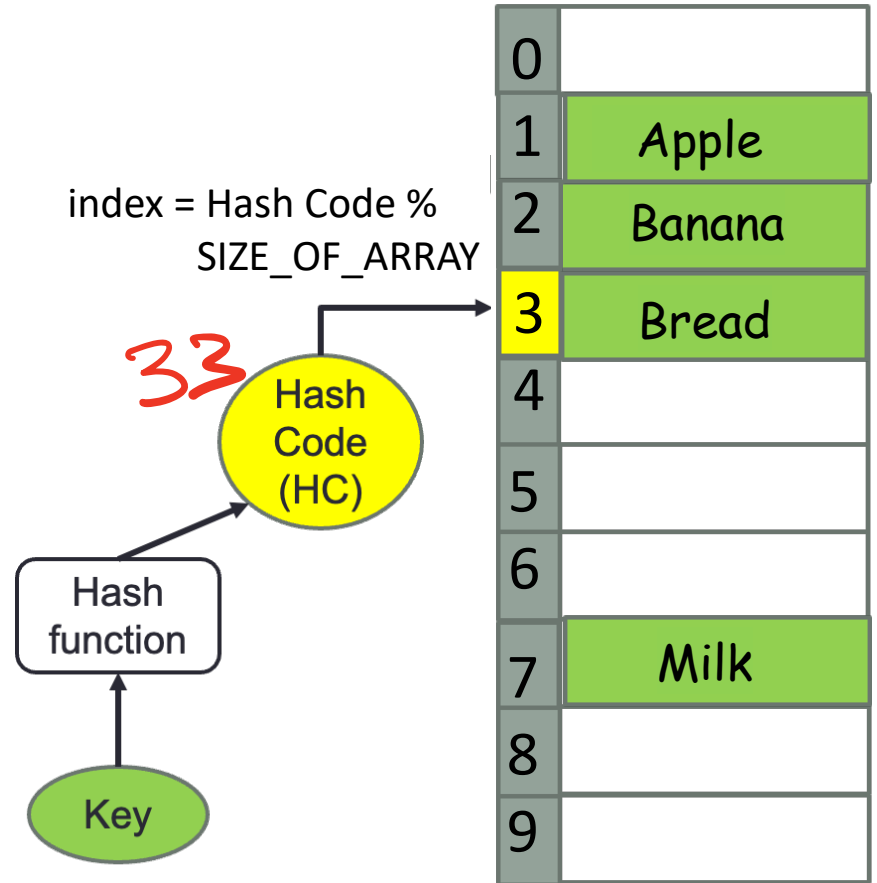
# Hash table is just an array with positions indexed by keys

```
unordered_set<string> groceries;  
groceries.insert("Banana");  
groceries.insert("Apple");  
groceries.insert("Milk");
```



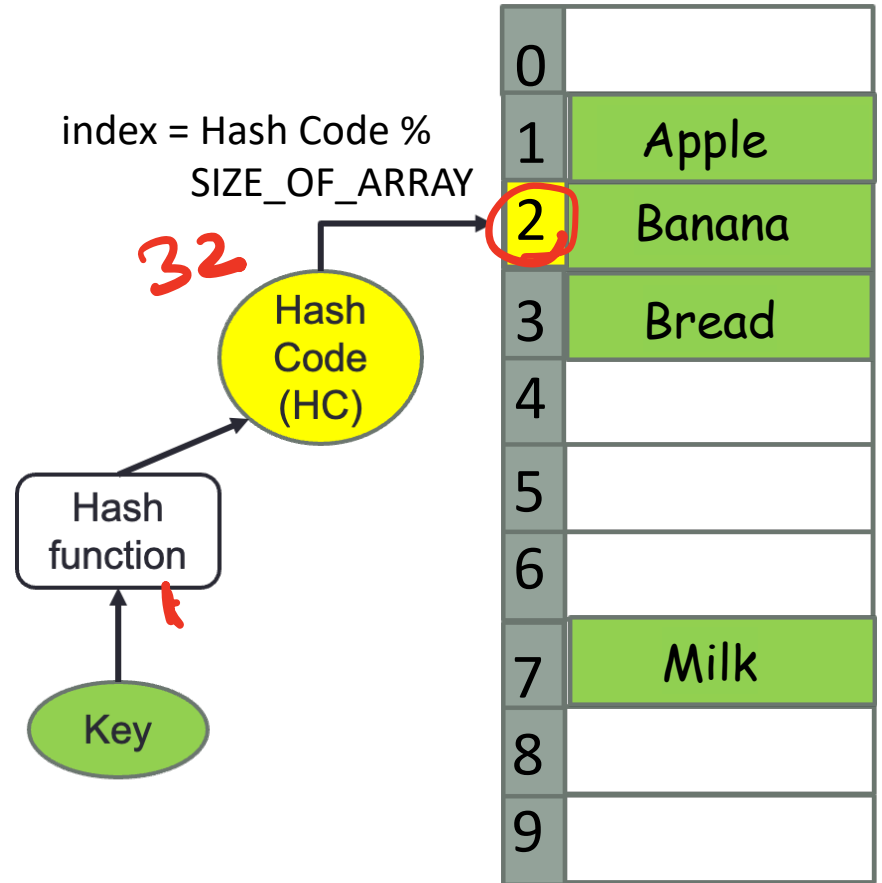
# Hash table is an array with positions indexed by keys

```
unordered_set<string> groceries;  
groceries.insert("Banana");  
groceries.insert("Apple");  
groceries.insert("Milk");  
groceries.insert("Bread");
```



# Hash table is an array with positions indexed by keys

```
unordered_set<string> groceries;  
groceries.insert("Banana");  
groceries.insert("Apple");  
groceries.insert("Milk");  
groceries.insert("Bread");  
groceries.find("Banana");
```

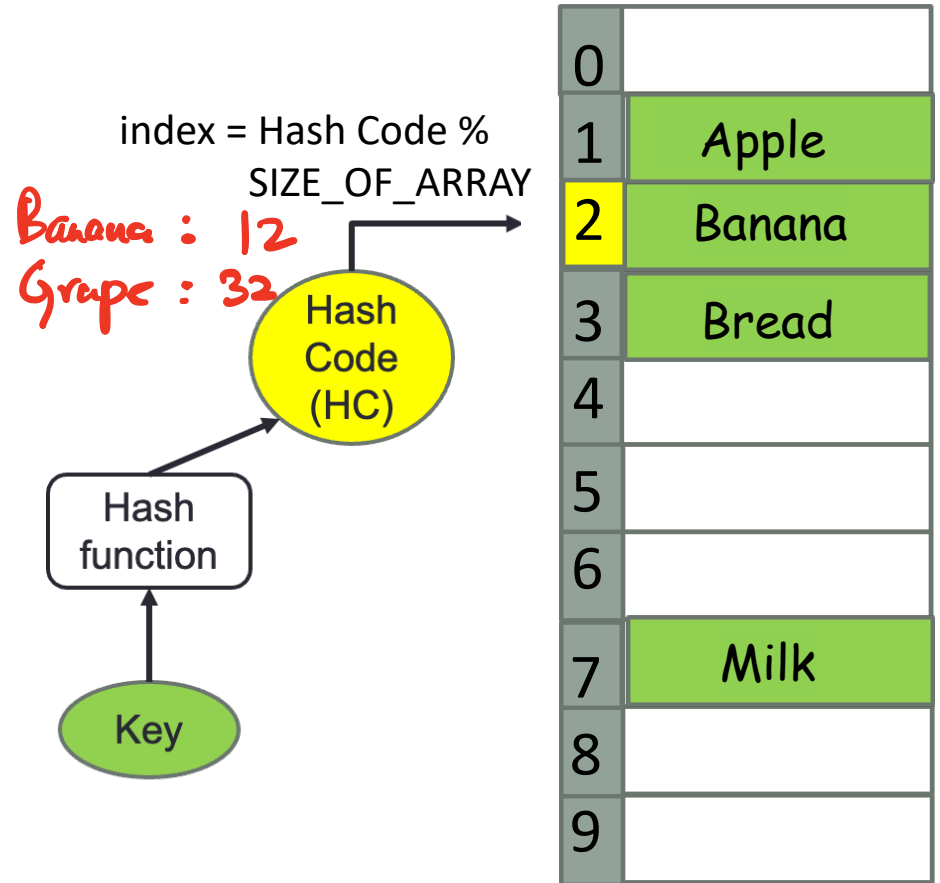


# Hash table is an array with positions indexed by keys

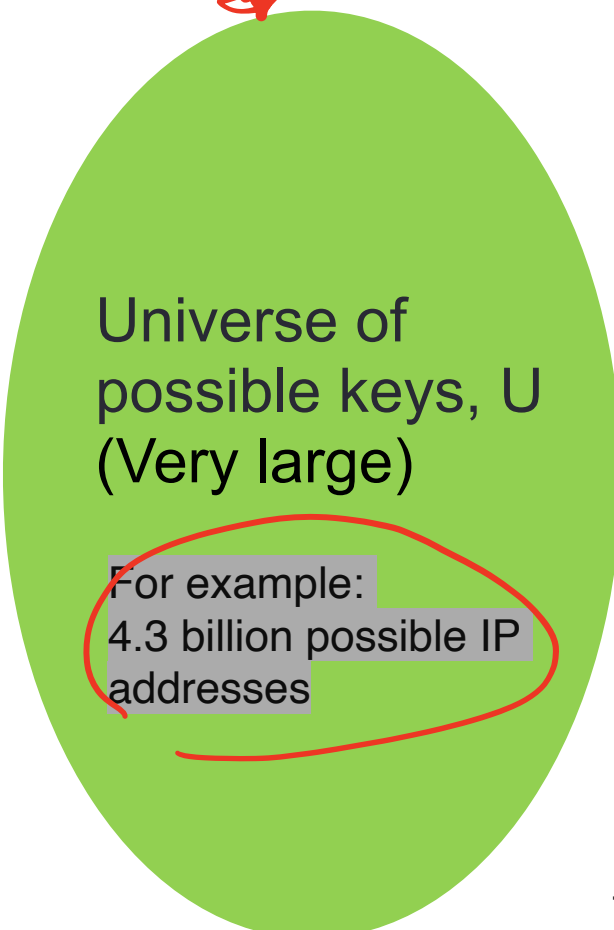
```
unordered_set<string> groceries;  
groceries.insert("Banana");  
groceries.insert("Apple");  
groceries.insert("Milk");  
groceries.insert("Bread");  
groceries.insert("Grape");
```

Suppose our hash function outputs the hash code 12 for "Banana" and 32 for "Grape", what happens when we try to insert "Grape"?

- A. "Grape" cannot be inserted
- B. "Banana" should be replaced by "Grape"
- C. Both "Banana" and "Grape" map to the same index, resulting in a collision




# Setup for hashing



Universe of possible keys,  $U$   
(Very large)

For example:  
4.3 billion possible IP  
addresses

Keep track of evolving set  $S$  whose size is much less than the universe of all possible keys



0	
1	
2	Key
3	
4	
5	
6	
7	
8	
9	

For example, size of a blacklist on an internet router is typically a few hundred to tens of thousands of entries

# Design challenges

- ✗ Deciding on collision resolution strategy
- ✗ Deciding the size of hash table
- Deciding the hash function

load factor

# of keys stored

# of indices in the table

Universe of possible keys,  $U$   
(Very large)

For example:  
4.3 billion possible IP addresses

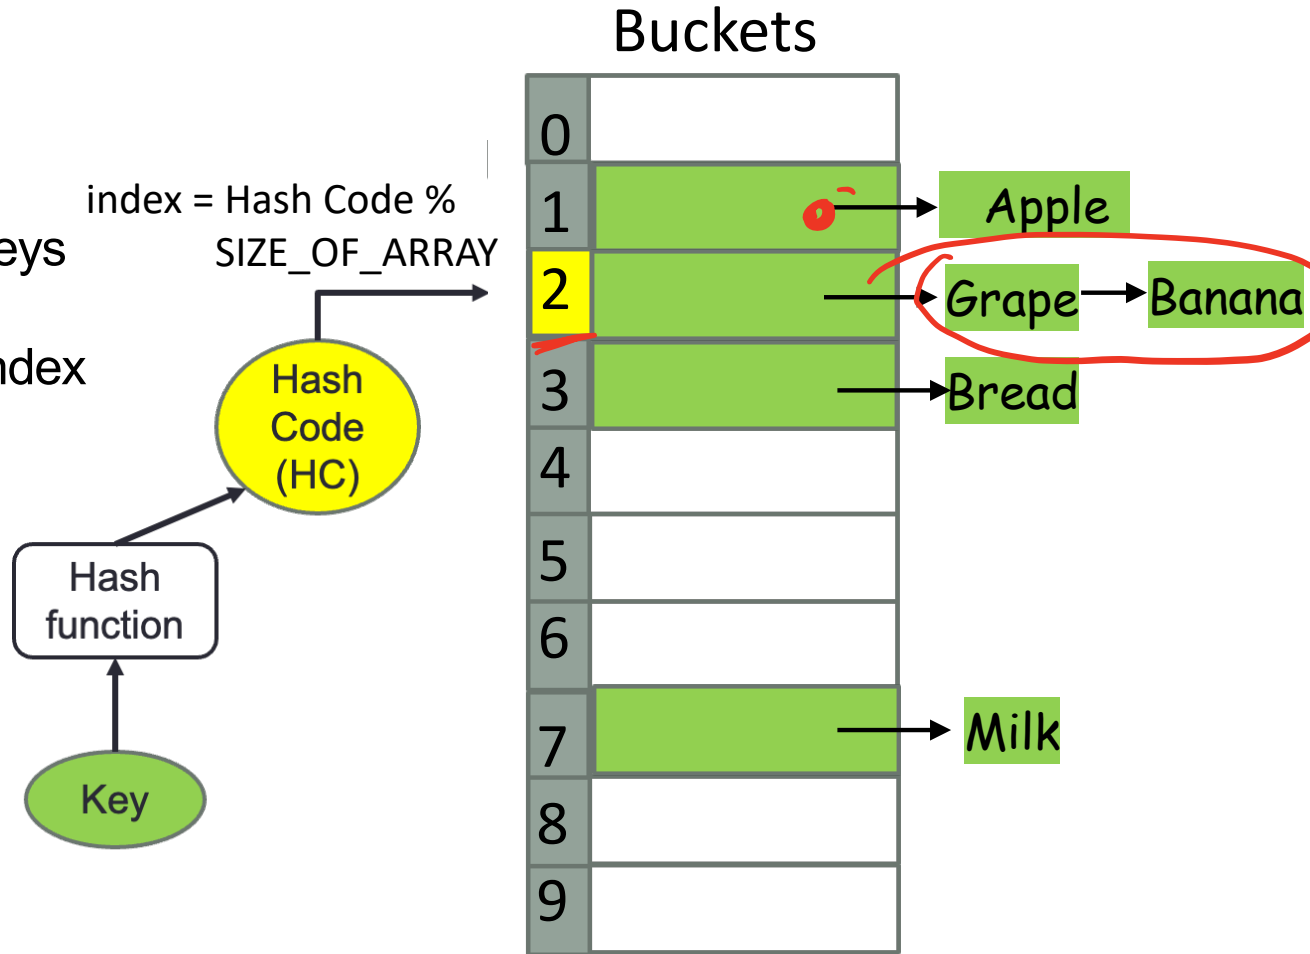
Keep track of evolving set  $S$  whose size is much less than the universe of all possible keys

0	
1	
2	Key
3	
4	
5	
6	
7	
8	
9	

For example, size of a blacklist on an internet router is typically a few hundred to tens of thousands of entries

# (Refined) Logical model of a hash table

- Keys stored in buckets (array)
- Array positions are indexed by keys
- Multiple keys map to the same index



# Two sum problem

Given an unsorted array (A) of N unique integers between 0 and 1000,000, find the pair of elements that sum to a given number T (assume a unique solution exists)

## Method 1: Exhaustive search

```
procedure twosum(A: array of integers of size N, T: target sum)
  pair<int, int> result
  for (i = 0; i < N; i++){
    for (j = i+1; j < N; j++){
      if (A[i] + A[j] == T)
        result = (A[i], A[j]);
    }
  }
  return result {pairs of values that add to T}
```

$O(N^2)$

What is the worst case run time of this method?

- A.  $O(N)$
- B.  $(N \log N)$
- C.  $O(N^2)$



# Two sum problem

Given an unsorted array (A) of N unique integers between 0 and 1000,000, find all pairs of elements that sum to a given number T (assume unique solution exists)

## Method 2: Sort, then search

```
procedure twosum(A: array of integers of size N, T: target sum)
  pair<int, int> result
  sort(A) O(N log N)
  for (i = 0; i < N; i++){
    if(binary_search(T - A[i])) O(log N)
      result = (A[i], T - A[i])
  }
  return result {pairs values that add to T}
```

What is the worst case run time of this method?

- A.  $O(N)$
- B.  $(N \log N)$
- C.  $O(N^2)$

# Two sum problem

Given an unsorted array (A) of N unique integers between 0 and 1000,000, find all pairs of elements that sum to a given number T

## Method 3: Use hash tables

```
procedure twosum(A: array of integers of size N, T: target sum)
  pair<int, int> result
```

*See code written in class*

```
return result {pairs of values add to T}
```

Discuss a method that uses hash tables. What is the worst case run time of this method?

- A.  $O(N)$
- B.  $(N \log N)$
- C.  $O(N^2)$

Suppose you have a hash table that can hold 100 elements. It currently stores 9 elements (in 9 different locations in the hash table). What is the probability that your next insert will cause a collision?

**Assume a hash function that maps keys to slots with equal likelihood**

- A. 0
- B.  $9/100$
- C.  $50/100$
- D.  $74/100$
- E. 1

Suppose you have a hash table that can hold 100 elements. It currently stores 30 elements (in one of 30 possible different locations in the hash table). What is the probability that your next two inserts will cause at least one collision? Assume a hash function that maps keys to slots with equal likelihood.

- A. 0
- B. 9/100
- C. 50/100
- D. 74/100
- E. 1

$$\begin{aligned} & \Pr(\text{collision}) \\ &= 1 - \Pr(\text{no collision}) \\ &= 1 - \Pr(\text{no collision first insert}) \wedge \Pr(\text{no collision in second insert}) \\ &= 1 - \frac{70}{100} * \frac{69}{100} \end{aligned}$$

Consider a party with  $n$  people, each person's birthday is on any day of the year with equal likelihood. How large does  $n$  need to be before there is a 50% chance that at least two people have the same birthday.

- A. 27
- B. 50
- C. 60
- D. 200
- E. 366

Suppose there are 365 slots in the hash table:  $M=365$

What is the probability that there will be a collision when inserting  $N$  keys?

For  $N = 10$ ,  $\text{prob}_{N,M}(\text{collision}) = 12\%$

For  $N = 20$ ,  $\text{prob}_{N,M}(\text{collision}) = 41\%$

For  $N = 30$ ,  $\text{prob}_{N,M}(\text{collision}) = 71\%$

For  $N = 40$ ,  $\text{prob}_{N,M}(\text{collision}) = 89\%$

For  $N = 50$ ,  $\text{prob}_{N,M}(\text{collision}) = 97\%$

For  $N = 60$ ,  $\text{prob}_{N,M}(\text{collision}) = 99+\%$

So, among 60 randomly selected people, it is almost certain that at least one pair of them have the same birthday

On average one pair of people will share a birthday in a group of about

$\sqrt{2 * 365} = 27$  people

In general: collisions are likely to happen, unless the hash table is quite sparsely filled

## Probability of collision

If a hash table has  $M$  slots and  $N$  keys, assuming your hashing function map keys to each slot with equal likelihood, the probability of at least one collision is

$$P_{N,M}(\text{collision}) = 1 - \prod_{i=1}^N \frac{(M - (i - 1))}{M}$$

# Hash table design (will be covered in depth in 130A)

- Deciding on collision resolution strategy
- Deciding the size of hash table
- Deciding the hash function

Universe of possible keys,  $U$   
(Very large)

For example:  
4.3 billion possible IP  
addresses

Keep track of evolving set  $S$   
whose size is much less than the  
universe of all possible keys

0	
1	
2	Key
3	
4	
5	
6	
7	
8	
9	

For example, size of a blacklist on an internet router is typically a few hundred to tens of thousands of entries