

Chapter 16

Hashing

OLD CS 473: Fundamental Algorithms, Spring 2015

March 17, 2015

16.1 Hash Tables

16.2 Introduction

16.2.0.1 Dictionary Data Structure

- (A) \mathcal{U} : universe of keys with total order: numbers, strings, etc.
- (B) Data structure to store a subset $S \subseteq \mathcal{U}$
- (C) **Operations:**
 - (A) **Search/lookup:** given $x \in \mathcal{U}$ is $x \in S$?
 - (B) **Insert:** given $x \notin S$ add x to S .
 - (C) **Delete:** given $x \in S$ delete x from S
- (D) **Static** structure: S given in advance or changes very infrequently, main operations are lookups.
- (E) **Dynamic** structure: S changes rapidly so inserts and deletes as important as lookups.


16.2.0.2 Dictionary Data Structures

Common solutions:

- (A) Static:
 - (A) Store S as a *sorted* array
 - (B) **Lookup:** Binary search in $O(\log |S|)$ time (comparisons)
- (B) Dynamic:
 - (A) Store S in a *balanced* binary search tree
 - (B) Lookup, Insert, Delete in $O(\log |S|)$ time (comparisons)

16.2.0.3 Dictionary Data Structures II

- (A) **Question:** “Should Tables be Sorted?”



(also title of famous paper by Turing award winner Andy Yao)

(B) Hashing is a widely used & powerful technique for dictionaries.

(C) **Motivation:**

(A) Universe \mathcal{U} may not be (naturally) totally ordered.

(B) Keys correspond to large objects (images, graphs etc) for which comparisons are very expensive.

(C) Want to improve “average” performance of lookups to $O(1)$ even at cost of extra space or errors with small probability: many applications for fast lookups in networking, security, etc.

16.2.0.4 Hashing and Hash Tables

(A) Hash Table data structure:

(A) A (hash) table/array T of size m (the table *size*).

(B) A hash function $h : \mathcal{U} \rightarrow \{0, \dots, m - 1\}$.

(C) Item $x \in \mathcal{U}$ hashes to slot $h(x)$ in T .

(B) Given $S \subseteq \mathcal{U}$. How do we store S and how do we do lookups?

(C) ...

Ideal situation:

(A) Each element $x \in S$ hashes to a distinct slot in T . Store x in slot $h(x)$

(B) **Lookup:** Given $y \in \mathcal{U}$ check if $T[h(y)] = y$. $O(1)$ time!

(D) Collisions unavoidable. Several different techniques to handle them.

16.2.0.5 Handling Collisions: Chaining

(A) **Collision:** $h(x) = h(y)$ for some $x \neq y$.

(B) **Chaining** to handle collisions:

(A) For each slot i store all items hashed to slot i in a linked list. $T[i]$ points to the linked list

(B) **Lookup:** to find if $y \in \mathcal{U}$ is in T , check the linked list at $T[h(y)]$. Time proportion to size of linked list.

(C) This is also known as **Open hashing**.

16.2.0.6 Handling Collisions

Several other techniques:

(A) Open addressing.

Every element has a list of places it can be (in certain order). Check in this order.

- (B) ...
- (C) Cuckoo hashing.
Every value has two possible locations. When inserting, insert in one of the locations, otherwise, kick stored value to its other location. Repeat till stable. if no stability then rebuild table.
- (D) Others.

16.2.0.7 Understanding Hashing

- (A) Does hashing give $O(1)$ time per operation for dictionaries?
- (B) **Questions:**
 - (A) Complexity of evaluating h on a given element?
 - (B) Relative sizes of the universe \mathcal{U} and the set to be stored S .
 - (C) Size of table relative to size of S .
 - (D) Worst-case vs average-case vs randomized (expected) time?
 - (E) How do we choose h ?

16.2.0.8 Understanding Hashing

- (A) Considerations:
 - (A) Complexity of evaluating h on a given element? Should be small.
 - (B) Relative sizes of the universe \mathcal{U} and the set to be stored S : typically $|\mathcal{U}| \gg |S|$.
 - (C) Size of table relative to size of S . The **load factor** of T is the ratio n/t where $n = |S|$ and $m = |T|$. Typically n/t is a small constant smaller than 1.
Also known as the **fill factor**.
- (B) Main and interrelated questions:
 - (A) Worst-case vs average-case vs randomized (expected) time?
 - (B) How do we choose h ?

16.2.0.9 Single hash function

- (A) \mathcal{U} : universe (very large).
- (B) Assume $N = |\mathcal{U}| \gg m$ where m is size of table T . In particular assume $N \geq m^2$ (very conservative).
- (C) Fix hash function $h : \mathcal{U} \rightarrow \{0, \dots, m - 1\}$.
- (D) N items hashed to m slots. By pigeon hole principle there is some $i \in \{0, \dots, m - 1\}$ such that $N/m \geq m$ elements of \mathcal{U} get hashed to i (!).
- (E) Implies that there is a set $S \subseteq \mathcal{U}$ where $|S| = m$ such that all of S hashes to same slot.
Oops.

Lesson: For every hash function there is a very bad set. Bad set. Bad.

16.2.0.10 Picking a hash function

- (A) How to pick functions?
 - (A) Hash function are often chosen in an ad hoc fashion. Implicit assumption is that input behaves well.

- (B) Theory and sound practice suggests that a hash function should be chosen properly with guarantees on its behavior.
- (B) Parameters: $N = |\mathcal{U}|$, $m = |T|$, $n = |S|$
 - (A) \mathcal{H} is a **family** of hash functions: each function $h \in \mathcal{H}$ should be efficient to evaluate (that is, to compute $h(x)$).
 - (B) h is chosen *randomly* from \mathcal{H} (typically uniformly at random). Implicitly assumes that \mathcal{H} allows an efficient sampling.
 - (C) Randomized guarantee: should have the property that for any *fixed* set $S \subseteq \mathcal{U}$ of size m the expected number of collisions for a function chosen from \mathcal{H} should be “small”. Here the expectation is over the randomness in choice of h .

16.2.0.11 Picking a hash function II

- (A) **Question:** Why not let \mathcal{H} be the set of *all* functions from \mathcal{U} to $\{0, 1, \dots, m - 1\}$?
- (B) (A) Too many functions! A random function has high complexity!
 # of functions: $M = m^{|\mathcal{U}|}$.
 Bits to encode such a function $\approx \log M = |\mathcal{U}| \log m$.
- (C) **Question:** Are there good and compact families \mathcal{H} ?
 (A) Yes... But what it means for \mathcal{H} to be good and compact.

16.3 Universal Hashing

16.3.0.12 Uniform hashing

Question: What are good properties of \mathcal{H} in distributing data?

- (A) Consider any element $x \in \mathcal{U}$. Then if $h \in \mathcal{H}$ is picked randomly then x should go into a random slot in T . In other words $\Pr[h(x) = i] = 1/m$ for every $0 \leq i < m$.
- (B) Consider any two distinct elements $x, y \in \mathcal{U}$. Then if $h \in \mathcal{H}$ is picked randomly then the probability of a collision between x and y should be at most $1/m$. In other words $\Pr[h(x) = h(y)] = 1/m$ (cannot be smaller).
- (C) Second property is stronger than the first and the crucial issue.

Definition 16.3.1. A family hash function \mathcal{H} is **2-universal** if for all distinct $x, y \in \mathcal{U}$, $\Pr[h(x) = h(y)] = 1/m$ where m is the table size.

Note: The set of all hash functions satisfies stronger properties!

16.3.0.13 Analyzing Uniform Hashing

- (A) T is hash table of size m .
- (B) $S \subseteq \mathcal{U}$ is a **fixed** set of size $\leq m$.
- (C) h is chosen randomly from a uniform hash family \mathcal{H} .
- (D) x is a *fixed* element of \mathcal{U} . Assume for simplicity that $x \notin S$.

Question: What is the *expected* time to look up x in T using h assuming chaining used to resolve collisions?

16.3.0.14 Analyzing Uniform Hashing

Question: What is the *expected* time to look up x in T using h assuming chaining used to resolve collisions?

- (A) The time to look up x is the size of the list at $T[h(x)]$: same as the number of elements in S that collide with x under h .
- (B) Let $\ell(x)$ be this number. We want $E[\ell(x)]$
- (C) For $y \in S$ let A_y be the event that x, y collide and D_y be the corresponding indicator variable.

16.3.1 Analyzing Uniform Hashing

16.3.1.1 Continued...

Number of elements colliding with x : $\ell(x) = \sum_{y \in S} D_y$.

$$\begin{aligned} \Rightarrow E[\ell(x)] &= \sum_{y \in S} E[D_y] && \text{linearity of expectation} \\ &= \sum_{y \in S} Pr[h(x) = h(y)] \\ &= \sum_{y \in S} \frac{1}{m} && \text{since } \mathcal{H} \text{ is a uniform hash family} \\ &= |S|/m \\ &\leq 1 && \text{if } |S| \leq m \end{aligned}$$

16.3.1.2 Analyzing Uniform Hashing

- (A) **Question:** What is the *expected* time to look up x in T using h assuming chaining used to resolve collisions?
- (B) **Answer:** $O(n/m)$.
- (C) Comments:
 - (A) $O(1)$ expected time also holds for insertion.
 - (B) Analysis assumes static set S but holds as long as S is a set formed with at most $O(m)$ insertions and deletions.
 - (C) **Worst-case:** look up time can be large! How large? $\Omega(\log n / \log \log n)$
[Lower bound holds even under stronger assumptions.]

16.3.2 Rehashing, amortization and...

16.3.2.1 ... making the hash table dynamic

Previous analysis assumed fixed S of size $\simeq m$.

Question: What happens as items are inserted and deleted?

- (A) If $|S|$ grows to more than cm for some constant c then hash table performance clearly degrades.
- (B) If $|S|$ stays around $\simeq m$ but incurs many insertions and deletions then the initial random hash function is no longer random enough!

Solution: Rebuild hash table periodically!

- (A) Choose a new table size based on current number of elements in table.
- (B) Choose a new random hash function and rehash the elements.
- (C) Discard old table and hash function.

Question: When to rebuild? How expensive?

16.3.2.2 Rebuilding the hash table

- (A) Start with table size m where m is some estimate of $|S|$ (can be some large constant).
- (B) If $|S|$ grows to more than twice current table size, build new hash table (choose a new random hash function) with double the current number of elements. Can also use similar trick if table size falls below quarter the size.
- (C) If $|S|$ stays roughly the same but more than $c|S|$ operations on table for some chosen constant c (say 10), rebuild.

The *amortize* cost of rebuilding to previously performed operations. Rebuilding ensures $O(1)$ expected analysis holds even when S changes. Hence $O(1)$ expected look up/insert/delete time *dynamic* data dictionary data structure!

16.3.2.3 Some math required...

Lemma 16.3.2. *Let p be a prime number,*

x : an integer number in $\{1, \dots, p-1\}$.

\implies There exists a unique y s.t. $xy = 1 \pmod p$.

In other words: For every element there is a unique inverse.

$\implies \mathbb{Z}_p = \{0, 1, \dots, p-1\}$ when working module p is a field.

16.3.2.4 Proof of lemma

Claim 16.3.3. *Let p be a prime number. For any $\alpha, \beta, i \in \{1, \dots, p-1\}$ s.t. $\alpha \neq \beta$, we have that $\alpha i \neq \beta i \pmod p$.*

Proof: Assume for the sake of contradiction $\alpha i = \beta i \pmod p$. Then

$$\begin{aligned}
 i(\alpha - \beta) &= 0 \pmod p \\
 \implies p &\text{ divides } i(\alpha - \beta) \\
 \implies p &\text{ divides } \alpha - \beta \\
 \implies \alpha - \beta &= 0 \\
 \implies \alpha &= \beta.
 \end{aligned}$$

And that is a contradiction. ■

16.3.3 Proof of lemma...

16.3.3.1 Uniqueness.

Lemma 16.3.4. *Let p be a prime number,
 x : an integer number in $\{1, \dots, p-1\}$.
 \implies There exists a unique y s.t. $xy = 1 \pmod p$.*

Proof: Assume the lemma is false and there are two distinct numbers $y, z \in \{1, \dots, p-1\}$ such that

$$xy = 1 \pmod p \quad \text{and} \quad xz = 1 \pmod p.$$

But this contradicts the above claim (set $i = x$, $\alpha = y$ and $\beta = z$). ■

16.3.4 Proof of lemma...

16.3.4.1 Existence

Proof: By claim, for any $\alpha \in \{1, \dots, p-1\}$ we have that $\{\alpha * 1 \pmod p, \alpha * 2 \pmod p, \dots, \alpha * (p-1) \pmod p\} = \{1, 2, \dots, p-1\}$.

\implies There exists a number $y \in \{1, \dots, p-1\}$ such that $\alpha y = 1 \pmod p$. ■

16.3.4.2 Constructing Universal Hash Families

Parameters: $N = |\mathcal{U}|$, $m = |T|$, $n = |S|$

(A) Choose a **prime** number $p \geq N$. $\mathbb{Z}_p = \{0, 1, \dots, p-1\}$ is a field.

(B) For $a, b \in \mathbb{Z}_p$, $a \neq 0$, define the hash function $h_{a,b}$ as $h_{a,b}(x) = ((ax + b) \pmod p) \pmod m$.

(C) Let $\mathcal{H} = \{h_{a,b} \mid a, b \in \mathbb{Z}_p, a \neq 0\}$. Note that $|\mathcal{H}| = p(p-1)$.

Theorem 16.3.5. \mathcal{H} is a 2-universal hash family.

Comments:


(A) Hash family is of small size, easy to sample from.

(B) Easy to store a hash function (a, b have to be stored) and evaluate it.

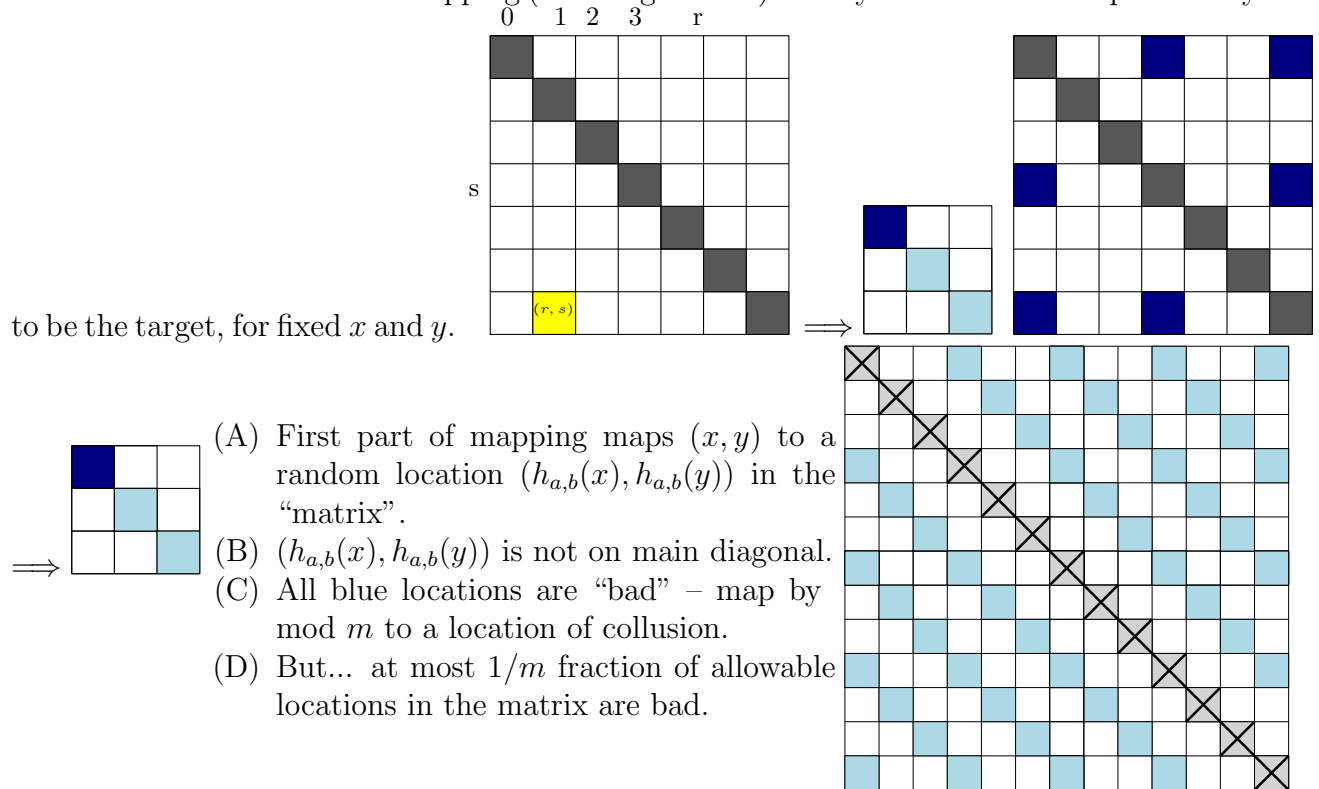
16.3.4.3 What the is going on?

$$h_{a,b}(x) = ((ax + b) \pmod p) \pmod m$$

First map $x \neq y$ to $r = h(x)$ and $s = h(y)$.



This is a random uniform mapping (choosing a and b) – every cell has the same probability



16.3.4.4 Constructing Universal Hash Families

Theorem 16.3.6. \mathcal{H} is a (2)-universal hash family.

Proof: Fix $x, y \in \mathcal{U}$. What is the probability they will collide if h is picked randomly from \mathcal{H} ?

- (A) Let a, b be bad for x, y if $h_{a,b}(x) = h_{a,b}(y)$.
- (B) **Claim:** Number of bad pairs is at most $p(p-1)/m$.
- (C) Total number of hash functions is $p(p-1)$ and hence probability of a collision is $\leq 1/m$. ■

16.3.4.5 Some Lemmas

Lemma 16.3.7. If $x \neq y$ then for any $a, b \in \mathbb{Z}_p$ such that $a \neq 0$, we have

$$ax + b \pmod p \neq ay + b \pmod p.$$

Proof: If $ax + b \pmod p = ay + b \pmod p$ then $a(x-y) \pmod p = 0$ and $a \neq 0$ and $(x-y) \neq 0$. However, a and $(x-y)$ cannot divide p since p is prime and $a < p$ and $(x-y) < p$. ■

16.3.4.6 Some Lemmas

Lemma 16.3.8. *If $x \neq y$ then for each (r, s) such that $r \neq s$ and $0 \leq r, s \leq p-1$ there is exactly **one** a, b such that*

$$ax + b \pmod p = r \text{ and } ay + b \pmod p = s$$

Proof: Solve the two equations:

$$ax + b = r \pmod p \quad \text{and} \quad ay + b = s \pmod p$$

We get $a = \frac{r-s}{x-y} \pmod p$ and $b = r - ax \pmod p$. ■

16.3.4.7 Understanding the hashing

Once we fix a and b , and we are given a value x , we compute the hash value of x in two stages:

(A) **Compute:** $r \leftarrow (ax + b) \pmod p$.

(B) **Fold:** $r' \leftarrow r \pmod m$

Collision...

Given two values x and y they might collide because of either steps.

Lemma 16.3.9. *# not equal pairs of $\mathbb{Z}_p \times \mathbb{Z}_p$ that are folded to the same number is $p(p-1)/m$.*

16.3.4.8 Folding numbers

Lemma 16.3.10. *# not equal pairs of $\mathbb{Z}_p \times \mathbb{Z}_p$ that are folded to the same number is $p(p-1)/m$.*

Proof: Consider a pair $(x, y) \in \{0, 1, \dots, p-1\}^2$ s.t. $x \neq y$. Fix x :

(A) There are $\lceil p/m \rceil$ values of y that fold into x . That is

$$x \pmod m = y \pmod m.$$

(B) One of them is when $x = y$.

(C) \implies # of colliding pairs $(\lceil p/m \rceil - 1)p \leq (p-1)p/m$ ■

16.3.5 Proof of Claim

16.3.5.1 # of bad pairs is $p(p-1)/m$

Proof: Let $a, b \in \mathbb{Z}_p$ such that $a \neq 0$ and $h_{a,b}(x) = h_{a,b}(y)$.

(A) Let $ax + b \pmod p = r$ and $ay + b \pmod p = s$.

(B) Collision if and only if $r = s \pmod m$.

- (C) (Folding error): Number of pairs (r, s) such that $r \neq s$ and $0 \leq r, s \leq p - 1$ and $r = s \pmod m$ is $p(p - 1)/m$.
- (D) From previous lemma for each bad pair (a, b) there is a unique pair (r, s) such that $r = s \pmod m$. Hence total number of bad pairs is $p(p - 1)/m$. ■

Prob of x and y to collide: $\frac{\# \text{ bad pairs}}{\# \text{ pairs}} = \frac{p(p-1)/m}{p(p-1)} = \frac{1}{m}$.

16.3.5.2 Perfect Hashing

- (A) **Question:** Can we make look up time $O(1)$ in worst case?
- (B) Yes, for static dictionaries but then space usage is $O(m)$ only in expectation.

16.3.5.3 Practical Issues

Hashing used typically for integers, vectors, strings etc.

- Universal hashing is defined for integers. To implement for other objects need to map objects in some fashion to integers (via representation)
- Practical methods for various important cases such as vectors, strings are studied extensively. See http://en.wikipedia.org/wiki/Universal_hashing for some pointers.
- Recent important paper bridging theory and practice of hashing. “The power of simple tabulation hashing” by Mikkel Thorup and Mihai Patrascu, 2011. See http://en.wikipedia.org/wiki/Tabulation_hashing

16.3.5.4 Bloom Filters

- (A) **Hashing:**
- (A) To insert x in dictionary store x in table in location $h(x)$
- (B) To lookup y in dictionary check contents of location $h(y)$
- (B) **Bloom Filter:** tradeoff space for false positives
- (A) Storing items in dictionary expensive in terms of memory, especially if items are unwieldy objects such a long strings, images, etc with *non-uniform* sizes.
- (B) To insert x in dictionary set *bit* to 1 in location $h(x)$ (initially all bits are set to 0)
- (C) To lookup y if bit in location $h(y)$ is 1 say yes, else no.

16.3.5.5 Bloom Filters

- (A) **Bloom Filter:** tradeoff space for false positives
- (A) To insert x in dictionary set *bit* to 1 in location $h(x)$ (initially all bits are set to 0)
- (B) To lookup y if bit in location $h(y)$ is 1 say yes, else no
- (C) No false negatives but false positives possible due to collisions
- (B) Reducing false positives:
- (A) Pick k hash functions h_1, h_2, \dots, h_k *independently*
- (B) To insert x for $1 \leq i \leq k$ set bit in location $h_i(x)$ in table i to 1

- (C) To lookup y compute $h_i(y)$ for $1 \leq i \leq k$ and say yes only if each bit in the corresponding location is 1, otherwise say no. If probability of false positive for one hash function is $\alpha < 1$ then with k independent hash function it is α^k .

16.3.5.6 Take away points

- (A) Hashing is a powerful and important technique for dictionaries. Many practical applications.
- (B) Randomization fundamental to understanding hashing.
- (C) Good and efficient hashing possible in theory and practice with proper definitions (universal, perfect, etc).
- (D) Related ideas of creating a compact fingerprint/sketch for objects is very powerful in theory and practice.
- (E) Many applications in practice.