## Lecture 8: Hashing

## Course Overview

This course covers several modules:

1. Review: dictionaries, chaining, simple uniform
2. Universal hashing
3. Perfect hashing

## Review

## Dictionary Problem

A dictionary is an Abstract Data Type (ADT) that maintains a set of items. Each item has a key. The dictionary supports the following operations:

- insert(item): add item to set
- delete(item): remove item from set
- search(key): return item with key if it exists

We assume that items have distinct keys (or that inserting new ones clobbers old ones).

This problem is easier than predecessor/successor problems solved in previous lecture (by van Emde Boas trees, or by AVL/2-3 trees/skip lists).

## Hashing from 6.006

Goal: $O(1)$ time per operation and $O(n)$ space complexity. Definitions:

- $u=$ number of keys over all possible items
- $n=$ number of keys/items currently in the table
- $m=$ number of slots in the table


Solution: hashing with chaining
Assuming simple uniform hashing,

$$
\operatorname{Pr}_{k_{1} \neq k_{2}}\left\{h\left(k_{1}\right)=h\left(k_{2}\right)\right\}=\frac{1}{m}
$$

we achieve $\Theta(1+\alpha)$ time per operation, where $\alpha=\frac{n}{m}$ is called load factor. The downside of the algorithm is that it requires assuming input keys are random, and it only works in average case, like basic quicksort. Today we are going to remove the unreasonable simple uniform hashing assumption.

## Etymology

The English 'hash' (1650s) means "cut into small pieces", which comes from the French 'hacher' which means "chop up", which comes from the Old French 'hache' which means "axe" (cf. English 'hatchet'). Alternatively, perhaps they come from Vulcan 'la'ash', which means "axe". (R.I.P. Leonard Nimoy.)

## Universal Hashing

The idea of universal hashing is listed as following:

- choose a random hash function $h$ from $\mathcal{H}$
- require $\mathcal{H}$ to be a universal hashing family such that

$$
\operatorname{Pr}_{h \in \mathcal{H}}\left\{h(k)=h\left(k^{\prime}\right)\right\} \leq \frac{1}{m} \text { for all } k \neq k^{\prime}
$$

- now we just assume $h$ is random, and make no assumption about input keys. (like Randomized Quicksort)

Theorem: For $n$ arbitrary distinct keys and random $h \in \mathcal{H}$, where $\mathcal{H}$ is a universal hashing family,

$$
E[\text { number of keys colliding in a slot }] \leq 1+\alpha \quad \text { where } \alpha=\frac{n}{m}
$$

Proof: Consider keys $k_{1}, k_{2}, \ldots, k_{n}$. Let $I_{i, j}=\left\{\begin{array}{ll}1 & \text { if } h\left(k_{i}\right)=h\left(k_{j}\right) \\ 0 & \text { otherwise }\end{array}\right.$. Then we have

$$
\begin{align*}
E\left[I_{i, j}\right] & =\operatorname{Pr}\left\{I_{i, j}=1\right\} \\
& =\operatorname{Pr}\left\{h\left(k_{i}\right)=h\left(k_{j}\right)\right\}  \tag{1}\\
& \leq \frac{1}{m} \text { for any } j \neq i
\end{align*}
$$

$E\left[\#\right.$ keys hashing to the same slot as $\left.k_{i}\right]=E\left[\sum_{j=1}^{n} I_{i, j}\right]$

$$
\begin{align*}
& =\sum_{j=1}^{n} E\left[I_{i, j}\right] \text { (linearity of expectation) } \\
& =\sum_{j \neq i} E\left[I_{i, j}\right]+E\left[I_{i, i}\right] \\
& \leq \frac{n}{m}+1 \tag{2}
\end{align*}
$$

From the above theorem, we know that Insert, Delete, and Search all take $O(1+\alpha)$ expected time. Here we give some examples of universal hash functions.

All hash functions: $\mathcal{H}=\{$ all hash functions $h:\{0,1, \ldots, u-1\} \rightarrow\{0,1, \ldots, m-$ $1\}\}$. Apparently, $\mathcal{H}$ is universal, but it is useless. On one hand, storing a single hashing function $h$ takes $\log \left(m^{u}\right)=u \log (m)$ bits $\gg n$ bits. On the other hand, we would need to precompute $u$ values, which takes $\Omega(u)$ time.

## Dot-product hash family:

Assumptions

- $m$ is a prime
- $u=m^{r}$ where $r$ is an integer

In real cases, we can always round up $m$ and $u$ to satisfy the above assumptions. Now let's view keys in base $m$ : $k=\left\langle k_{0}, k_{1}, \ldots, k_{r-1}\right\rangle$. For key $a=\left\langle a_{0}, a_{1}, a_{2}, \ldots, a_{r-1}\right\rangle$, define

$$
\begin{align*}
h_{a}(k) & =a \cdot k \bmod m(\text { dot product }) \\
& =\sum_{i=0}^{r-1} a_{i} k_{i} \bmod m \tag{3}
\end{align*}
$$

Then our hash family is $\mathcal{H}=\left\{h_{a} \mid a \in\{0,1, \ldots, u-1\}\right\}$
Storing $h_{a} \in \mathcal{H}$ requires just storing one key, which is $a$. In the word RAM model, manipulating $O(1)$ machine words takes $O(1)$ time and "objects of interest" (here, keys) fit into a machine word. Thus computing $h_{a}(k)$ takes $O(1)$ time.

Theorem: Dot-product hash family $\mathcal{H}$ is universal.

Proof: Take any two keys $k \neq k^{\prime}$. They must differ in some digits. Say $k_{d} \neq k_{d^{\prime}}$. Define not $d=\{0,1, \ldots, r-1\} \backslash\{d\}$. Now we have

$$
\begin{align*}
\operatorname{Pr}_{a}\left\{h_{a}(k)=h_{a}\left(k^{\prime}\right)\right\} & =\operatorname{Pr}_{a}\left\{\sum_{i=0}^{r-1} a_{i} k_{i}=\sum_{i=0}^{r-1} a_{i} k_{i}^{\prime}(\bmod m)\right\} \\
& =\operatorname{Pr}_{a}\left\{\sum_{i \neq d} a_{i} k_{i}+a_{d} k_{d}=\sum_{i \neq d} a_{i} k_{i}^{\prime}+a_{d} k_{d}^{\prime}(\bmod m)\right\} \\
& =\operatorname{Pr}_{a}\left\{\sum_{i \neq d} a_{i}\left(k_{i}-k_{i}^{\prime}\right)+a_{d}\left(k_{d}-k_{d}^{\prime}\right)=0(\bmod m)\right\} \\
& =\operatorname{Pr}_{a}\left\{a_{d}=-\left(k_{d}-k_{d}^{\prime}\right)^{-1} \sum_{i \neq d} a_{i}\left(k_{i}-k_{i}^{\prime}\right)(\bmod m)\right\}  \tag{4}\\
& \left(m \text { is prime } \Rightarrow \mathbb{Z}_{m} \text { has multiplicative inverses }\right) \\
& =\underset{a_{\text {not }}}{E}\left[\operatorname{Par}_{d}\left\{a_{d}=f\left(k, k^{\prime}, a_{\text {not } d}\right)\right\}\right] \\
& \left(=\sum_{x} \operatorname{Pr}\left\{a_{\text {not } d}=x\right\} \operatorname{Pr}_{a_{d}}\left\{a_{d}=f\left(k, k^{\prime}, x\right)\right\}\right) \\
& =\underset{a_{\text {not }}}{E}\left[\frac{1}{m}\right] \\
& =\frac{1}{m}
\end{align*}
$$

Another universal hash family from CLRS: Choose prime $p \geq u$ (once). Define $\left.h_{a b}(k)=[(a k+b) \bmod p)\right] \bmod m$. Let $\mathcal{H}=\left\{h_{a b} \mid a, b \in\{0,1, \ldots, u-1\}\right\}$.

## Perfect Hashing

Static dictionary problem: Given $n$ keys to store in table, only need to support $\operatorname{search}(k)$. No insertion or deletion will happen.

Perfect hashing: [Fredman, Komlós, Szemerédi 1984]

- polynomial build time with high probability (w.h.p.)
- $O(1)$ time for search in worst case
- $O(n)$ space in worst case

Idea: 2-level hashing


The algorithm contains the following two major steps:
Step 1: Pick $h_{1}:\{0,1, \ldots, u-1\} \rightarrow\{0,1, \ldots, m-1\}$ from a universal hash family for $m=\Theta(n)$ (e.g., nearby prime). Hash all items with chaining using $h_{1}$.

Step 2: For each slot $j \in\{0,1, \ldots, m-1\}$, let $l_{j}$ be the number of items in slot $j$. $l_{j}=\left|\left\{i \mid h\left(k_{i}\right)=j\right\}\right|$. Pick $h_{2, j}:\{0,1, \ldots, u-1\} \rightarrow\left\{0,1, \ldots, m_{j}\right\}$ from a universal hash family for $l_{j}^{2} \leq m_{j} \leq O\left(l_{j}^{2}\right)$ (e.g., nearby prime). Replace chain in slot $j$ with hashing-with-chaining using $h_{2, j}$.

The space complexity is $O\left(n+\sum_{j=0}^{m-1} l_{j}^{2}\right)$. In order to reduce it to $O(n)$, we need to add two more steps:

Step 1.5: If $\sum_{j=0}^{m-1} l_{j}^{2}>c n$ where $c$ is a chose constant, then redo Step 1.
Step 2.5: While $h_{2, j}\left(k_{i}\right)=h_{2, j}\left(k_{i}^{\prime}\right)$ for any $i \neq i^{\prime}, j$, repick $h_{2, j}$ and rehash those $l_{j}$.

The above two steps guarantee that there are no collisions at second level, and the space complexity is $O(n)$. As a result, search time is $O(1)$. Now let's look at the build time of the algorithm. Both Step 1 and Step 2 are $O(n)$. How about Step 1.5 and Step 2.5?

For Step 2.5,

$$
\begin{align*}
\underset{h_{2, j}}{\operatorname{Pr}}\left\{h_{2, j}\left(k_{i}\right)=h_{2, j}\left(k_{i}^{\prime}\right) \text { for some } i \neq i^{\prime}\right\} & \leq \sum_{i \neq i^{\prime}} \operatorname{Pr}_{2, j}\left\{h_{2, j}\left(k_{i}\right)=h_{2, j}\left(k_{i}^{\prime}\right)\right\} \text { (union bound) } \\
& \leq\binom{ l_{j}}{2} \cdot \frac{1}{l_{j}^{2}} \\
& <\frac{1}{2} \tag{5}
\end{align*}
$$

As a result, each trial is like a coin flip. If the outcome is "tail", we move to the next step. By Lecture 7 , we have $E[\#$ trials $] \leq 2$ and \#trials $=O(\log n)$ w.h.p. By a Chernoff bound, $l_{j}=O(\log n)$ w.h.p., so each trial takes $O(\log n)$ time. Because we have to do this for each $j$, the total time complexity is $O(\log n) \cdot O(\log n) \cdot O(n)=$ $O\left(n \log ^{2} n\right)$ w.h.p.

$$
\begin{align*}
& \text { For Step 1.5, we define } I_{i, i^{\prime}}= \begin{cases}1 & \text { if } h\left(k_{i}\right)=h\left(k_{i}^{\prime}\right) . \text { Then we have } \\
0 & \text { otherwise }\end{cases} \\
& \qquad \begin{aligned}
E\left[\sum_{j=0}^{m-1} l_{j}^{2}\right] & =E\left[\sum_{i=1}^{n} \sum_{i^{\prime}=1}^{n} I_{i, i^{\prime}}\right] \\
& =\sum_{i=1}^{n} \sum_{i^{\prime}=1}^{n} E\left[I_{i, i^{\prime}}\right] \text { (linearity of expectation) } \\
& \leq n+2\binom{n}{2} \cdot \frac{1}{m} \\
& =O(n) \text { because } m=\Theta(n)
\end{aligned}
\end{align*}
$$

By Markov inequality, we have

$$
\underset{h_{1}}{\operatorname{Pr}}\left\{\sum_{j=0}^{m-1} l_{j}^{2} \leq c n\right\} \leq \frac{\sum_{j=0}^{m-1} l_{j}^{2}}{c n} \leq \frac{1}{2}
$$

for a sufficiently large constant $c$. By Lecture 7 , we have $E[\#$ trials $] \leq 2$ and \#trials $=$ $O(\log n)$ w.h.p. As a result, Step 1 and Step 1.5 combined takes $O(n \log n)$ time w.h.p.

MIT OpenCourseWare
http://ocw.mit.edu

### 6.046J / 18.410J Design and Analysis of Algorithms

Spring 2015

For information about citing these materials or our Terms of Use, visit: http://ocw.mit.edu/terms.

