CSE373: Data Structures \& Algorithms Lecture 17: Hash Collisions

Kevin Quinn
Fall 2015

## Hash Tables: Review

- Aim for constant-time (i.e., O(1)) find, insert, and delete
- "On average" under some reasonable assumptions
- A hash table is an array of some fixed size
- But growable as we'll see


TableSize - 1

## Collision resolution

## Collision:

When two keys map to the same location in the hash table

We try to avoid it, but number-of-keys exceeds table size

So hash tables should support collision resolution

- Ideas?


## Separate Chaining

| 0 | 1 |
| :---: | :---: |
| 1 | 1 |
| 2 | 1 |
| 3 | 1 |
| 4 | 1 |
| 5 | 1 |
| 6 | 1 |
| 7 | 1 |
| 8 | 1 |
| 9 | 1 |

## Chaining:

All keys that map to the same table location are kept in a list (a.k.a. a "chain" or "bucket")

As easy as it sounds

## Example:

insert 10, 22, 107, 12, 42
with mod hashing
and TableSize $=10$

## Separate Chaining



Chaining:
All keys that map to the same table location are kept in a list (a.k.a. a "chain" or "bucket")

As easy as it sounds

Example:
insert 10, 22, 107, 12, 42
with mod hashing
and TableSize $=10$

## Separate Chaining



Chaining:
All keys that map to the same table location are kept in a list (a.k.a. a "chain" or "bucket")

As easy as it sounds

Example:
insert 10, 22, 107, 12, 42
with mod hashing
and TableSize $=10$

## Separate Chaining



Chaining:
All keys that map to the same table location are kept in a list (a.k.a. a "chain" or "bucket")

As easy as it sounds

Example:
insert 10, 22, 107, 12, 42
with mod hashing
and TableSize $=10$

## Separate Chaining



Chaining:
All keys that map to the same table location are kept in a list (a.k.a. a "chain" or "bucket")

As easy as it sounds

Example:
insert 10, 22, 107, 12, 42
with mod hashing
and TableSize $=10$

## Separate Chaining



Chaining:
All keys that map to the same table location are kept in a list (a.k.a. a "chain" or "bucket")

As easy as it sounds

Example:
insert 10, 22, 107, 12, 42
with mod hashing and TableSize $=10$

## Thoughts on chaining

- Worst-case time for find?
- Linear
- But only with really bad luck or bad hash function
- So not worth avoiding (e.g., with balanced trees at each bucket)
- Beyond asymptotic complexity, some "data-structure engineering" may be warranted
- Linked list vs. array vs. chunked list (lists should be short!)
- Move-to-front
- Maybe leave room for 1 element (or 2?) in the table itself, to optimize constant factors for the common case
- A time-space trade-off...


## Time vs. space (constant factors only here)



## More rigorous chaining analysis

Definition: The load factor, $\lambda$, of a hash table is

$$
\lambda=\frac{\mathrm{N}}{\text { TableSize }} \leftarrow \text { number of elements }
$$

Under chaining, the average number of elements per bucket is $\lambda$

So if some inserts are followed by random finds, then on average:

- Each "unsuccessful" find compares against $\lambda$ items

So we like to keep $\lambda$ fairly low (e.g., 1 or 1.5 or 2 ) for chaining

## Alternative: Use empty space in the table

- Another simple idea: If $h$ (key) is already full,
- try (h(key) + 1) \% TableSize. If full,
- try (h(key) + 2) \% TableSize. If full,
- try (h(key) + 3) \% TableSize. If full...
- Example: insert 38, 19, 8, 109, 10

| 0 | 1 |
| :---: | :---: |
| 1 | 1 |
| 2 | 1 |
| 3 | 1 |
| 4 | / |
| 5 | 1 |
| 6 | 1 |
| 7 | 1 |
| 8 | 38 |
| 9 | / |

## Alternative: Use empty space in the table

- Another simple idea: If $h$ (key) is already full,
- try (h(key) + 1) \% TableSize. If full,
- try (h(key) + 2) \% TableSize. If full,
- try (h(key) + 3) \% TableSize. If full...
- Example: insert 38, 19, 8, 109, 10

| 0 | / |
| :---: | :---: |
| 1 | / |
| 2 | / |
| 3 | / |
| 4 | 1 |
| 5 | 1 |
| 6 | 1 |
| 7 | 1 |
| 8 | 38 |
| 9 | 19 |

## Alternative: Use empty space in the table

- Another simple idea: If $h$ (key) is already full,
- try (h(key) + 1) \% TableSize. If full,
- try (h(key) + 2) \% TableSize. If full,
- try (h(key) + 3) \% TableSize. If full...
- Example: insert 38, 19, 8, 109, 10

| 0 | 8 |
| :---: | :---: |
| 1 | / |
| 2 | / |
| 3 | 1 |
| 4 | / |
| 5 | 1 |
| 6 | 1 |
| 7 | 1 |
| 8 | 38 |
| 9 | 19 |

## Alternative: Use empty space in the table

- Another simple idea: If $h$ (key) is already full,
- try (h(key) + 1) \% TableSize. If full,
- try (h (key) + 2) \% TableSize. If full,
- try (h(key) + 3) \% TableSize. If full...
- Example: insert 38, 19, 8, 109, 10

| 0 | 8 |
| :---: | :---: |
| 1 | 109 |
| 2 | / |
| 3 | / |
| 4 | / |
| 5 | / |
| 6 | / |
| 7 | / |
| 8 | 38 |
| 9 | 19 |

## Alternative: Use empty space in the table

- Another simple idea: If $h$ (key) is already full,
- try (h(key) + 1) \% TableSize. If full,
- try (h (key) + 2) \% TableSize. If full,
- try (h(key) + 3) \% TableSize. If full...
- Example: insert 38, 19, 8, 109, 10

| 0 | 8 |
| :---: | :---: |
| 1 | 109 |
| 2 | 10 |
| 3 | 1 |
| 4 | / |
| 5 | 1 |
| 6 | 1 |
| 7 | 1 |
| 8 | 38 |
| 9 | 19 |

## Open addressing

This is one example of open addressing
In general, open addressing means resolving collisions by trying a sequence of other positions in the table

Trying the next spot is called probing

- We just did linear probing
- $i^{\text {th }}$ probe was (h(key) + i) \% TableSize
- In general have some probe function $\mathbf{f}$ and use $h(k e y)+f(i) \%$ TableSize

Open addressing does poorly with high load factor $\lambda$

- So want larger tables
- Too many probes means no more $O(1)$


## Terminology

We and the book use the terms

- "chaining" or "separate chaining"
- "open addressing"

Very confusingly,

- "open hashing" is a synonym for "chaining"
- "closed hashing" is a synonym for "open addressing"
(If it makes you feel any better, most trees in CS grow upside-down () )



## Other operations

insert finds an open table position using a probe function

What about find?

- Must use same probe function to "retrace the trail" for the data
- Unsuccessful search when reach empty position

What about delete?

- Must use "lazy" deletion. Why?
- Marker indicates "no data here, but don't stop probing"
- Note: delete with chaining is plain-old list-remove


## (Primary) Clustering

It turns out linear probing is a bad idea, even though the probe function is quick to compute (which is a good thing)

Tends to produce clusters, which lead to long probing sequences

- Called primary clustering
- Saw this starting in our example



## Analysis of Linear Probing

- Trivial fact: For any $\lambda<1$, linear probing will find an empty slot
- It is "safe" in this sense: no infinite loop unless table is full
- Non-trivial facts we won't prove:

Average \# of probes given $\lambda$ (in the limit as TableSize $\rightarrow \infty$ )

- Unsuccessful search:

$$
\frac{1}{2}\left(1+\frac{1}{(1-\lambda)^{2}}\right)
$$

- Successful search:

$$
\frac{1}{2}\left(1+\frac{1}{(1-\lambda)}\right)
$$

- This is pretty bad: need to leave sufficient empty space in the table to get decent performance


## In a chart

- Linear-probing performance degrades rapidly as table gets full
- (Formula assumes "large table" but point remains)

- By comparison, chaining performance is linear in $\lambda$ and has no trouble with $\lambda>1$


## Quadratic probing

- We can avoid primary clustering by changing the probe function (h(key) $+f(i)$ ) \% TableSize
- A common technique is quadratic probing:

$$
f(i)=i^{2}
$$

- So probe sequence is:
- $0^{\text {th }}$ probe: $\mathrm{h}($ key $) ~ \% ~ T a b l e S i z e ~$
- $1^{\text {st }}$ probe: $(\mathrm{h}(\mathrm{key})+1) \%$ TableSize
- $2^{\text {nd }}$ probe: $(\mathrm{h}(\mathrm{key})+4) \%$ TableSize
- $3^{\text {rd }}$ probe: $(\mathrm{h}(\mathrm{key})+9) \%$ TableSize
- ...
- $i^{\text {th }}$ probe: $\left(\mathrm{h}(\right.$ key $\left.)+\mathrm{i}^{2}\right) \%$ TableSize
- Intuition: Probes quickly "leave the neighborhood"


## Quadratic Probing Example



TableSize=10 Insert:
89
18
49
58
79

## Quadratic Probing Example



TableSize=10 Insert:

## Quadratic Probing Example



TableSize=10 Insert:
89
18
49
58
79

## Quadratic Probing Example



TableSize=10 Insert:

## Quadratic Probing Example



TableSize=10 Insert:

## Quadratic Probing Example



TableSize=10 Insert:
89
18
49
58
79

## Another Quadratic Probing Example



TableSize $=7$

| Insert: |  |
| :--- | :--- |
| 76 | $(76 \% 7=6)$ |
| 40 | $(40 \% 7=5)$ |
| 48 | $(48 \% 7=6)$ |
| 5 | $(5 \% 7=5)$ |
| 55 | $(55 \% 7=6)$ |
| 47 | $(47 \% 7=5)$ |

## Another Quadratic Probing Example



TableSize $=7$

| Insert: |  |
| :--- | :--- |
| 76 | $(76 \% 7=6)$ |
| 40 | $(40 \% 7=5)$ |
| 48 | $(48 \% 7=6)$ |
| 5 | $(5 \% 7=5)$ |
| 55 | $(55 \% 7=6)$ |
| 47 | $(47 \% 7=5)$ |

## Another Quadratic Probing Example



TableSize $=7$

| Insert: |  |
| :--- | :--- |
| 76 | $(76 \% 7=6)$ |
| 40 | $(40 \% 7=5)$ |
| 48 | $(48 \% 7=6)$ |
| 5 | $(5 \% 7=5)$ |
| 55 | $(55 \% 7=6)$ |
| 47 | $(47 \% 7=5)$ |

## Another Quadratic Probing Example



TableSize $=7$

| Insert: |  |
| :--- | :--- |
| 76 | $(76 \% 7=6)$ |
| 40 | $(40 \% 7=5)$ |
| 48 | $(48 \% 7=6)$ |
| 5 | $(5 \% 7=5)$ |
| 55 | $(55 \% 7=6)$ |
| 47 | $(47 \% 7=5)$ |

## Another Quadratic Probing Example



TableSize $=7$

| Insert: |  |
| :--- | :--- |
| 76 | $(76 \% 7=6)$ |
| 40 | $(40 \% 7=5)$ |
| 48 | $(48 \% 7=6)$ |
| 5 | $(5 \% 7=5)$ |
| 55 | $(55 \% 7=6)$ |
| 47 | $(47 \% 7=5)$ |

## Another Quadratic Probing Example



TableSize $=7$

| Insert: |  |
| :--- | :--- |
| 76 | $(76 \% 7=6)$ |
| 40 | $(40 \% 7=5)$ |
| 48 | $(48 \% 7=6)$ |
| 5 | $(5 \% 7=5)$ |
| 55 | $(55 \% 7=6)$ |
| 47 | $(47 \% 7=5)$ |

## Another Quadratic Probing Example



TableSize = 7

| Insert: |  |
| :--- | :--- |
| 76 | $(76 \% 7=6)$ |
| 40 | $(40 \% 7=5)$ |
| 48 | $(48 \% 7=6)$ |
| 5 | $(5 \% 7=5)$ |
| 55 | $(55 \% 7=6)$ |
| 47 | $(47 \% 7=5)$ |

Doh!: For all $n,\left(\left(n *_{n}\right)+5\right) \% 7$ is 0,2 , 5 , or 6

- Excel shows takes "at least" 50 probes and a pattern
- Proof uses induction and $\left(n^{2}+5\right) \div 7=\left((n-7)^{2}+5\right) \% 7$
- In fact, for all $c$ and $\left.k,\left(n^{2}+c\right) \% k=(n-k)^{2}+c\right) \% k$


## From Bad News to Good News

- Bad news:
- Quadratic probing can cycle through the same full indices, never terminating despite table not being full
- Good news:
- If TableSize is prime and $\lambda<1 / 2$, then quadratic probing will find an empty slot in at most TableSize/2 probes
- So: If you keep $\lambda<1 / 2$ and TableSize is prime, no need to detect cycles
- Optional
- Also, slightly less detailed proof in textbook
- Key fact: For prime $\mathbf{T}$ and $0<\mathbf{i}, \mathbf{j}<\mathbf{T} / 2$ where $\mathbf{i} \neq \mathbf{j}$, $\left(\mathbf{k}+\mathbf{i}^{2}\right) \% \mathbf{T} \neq\left(\mathbf{k}+\mathbf{j}^{2}\right) \% \mathbf{T}$ (i.e., no index repeat)


## Clustering reconsidered

- Quadratic probing does not suffer from primary clustering: no problem with keys initially hashing to the same neighborhood
- But it's no help if keys initially hash to the same index
- Called secondary clustering
- Can avoid secondary clustering with a probe function that depends on the key: double hashing...


## Double hashing

Idea:

- Given two good hash functions $h$ and $g$, it is very unlikely that for some key, h (key) == g (key)
- So make the probe function $\mathrm{f}(\mathrm{i})=\mathrm{i}$ * g (key)

Probe sequence:

- $0^{\text {th }}$ probe: $\mathrm{h}(\mathrm{key}) ~ \% ~ T a b l e S i z e ~$
- $1^{\text {st }}$ probe: (h(key) $+\mathrm{g}(\mathrm{key})$ ) $\%$ TableSize
- $2^{\text {nd }}$ probe: (h(key) +2 *g(key)) $\%$ TableSize
- $3^{\text {rd }}$ probe: (h(key) + 3*g(key)) \% TableSize
- $i^{\text {th }}$ probe: (h (key) + i*g(key)) $\%$ TableSize

Detail: Make sure g(key) cannot be 0

## Double-hashing analysis

- Intuition: Because each probe is "jumping" by g(key) each time, we "leave the neighborhood" and "go different places from other initial collisions"
- But we could still have a problem like in quadratic probing where we are not "safe" (infinite loop despite room in table)
- It is known that this cannot happen in at least one case:
- h(key) = key \% p
- g(key) = q - (key \% q)
- $2<q<p$
- p and q are prime


## More double-hashing facts

- Assume "uniform hashing"
- Means probability of $g($ key 1$) \% p==g($ key 2$) \% p$ is 1/p
- Non-trivial facts we won't prove:

Average \# of probes given $\lambda$ (in the limit as TableSize $\rightarrow \infty$ )

- Unsuccessful search (intuitive):

$$
\frac{1}{1-\lambda}
$$

- Successful search (less intuitive):

$$
\frac{1}{\lambda} \log _{e}\left(\frac{1}{1-\lambda}\right)
$$

- Bottom line: unsuccessful bad (but not as bad as linear probing), but successful is not nearly as bad


## Charts



## Rehashing

- As with array-based stacks/queues/lists, if table gets too full, create a bigger table and copy everything
- With chaining, we get to decide what "too full" means
- Keep load factor reasonable (e.g., < 1 )?
- Consider average or max size of non-empty chains?
- For open addressing, half-full is a good rule of thumb
- New table size
- Twice-as-big is a good idea, except, uhm, that won't be prime!
- So go about twice-as-big
- Can have a list of prime numbers in your code since you won't grow more than 20-30 times

