Hash Tables – Double hashing

Carlos Moreno
cmoreno@uwaterloo.ca
EIT-4103

https://ece.uwaterloo.ca/~cmoreno/ece250

These slides, the course material, and course web site are based on work by Douglas W. Harder

Hash Tables – Double hashing

Standard reminder to set phones to silent/vibrate mode, please!

Hash Tables – Double hashing

- Today's class:
  - We'll look at one of the issues with linear probing, namely clustering
  - Discuss double hashing:
    - Use one hash function to determine the bin
    - A second hash function determines the jump size for the probing sequence.
  - Look at some practical issues and approaches to deal with these issues.

Hash Tables – Double hashing

- One important problem with linear probing is clustering — as collisions start to occur, then blocks of contiguous occupied bins (clusters) appear.

Hash Tables – Double hashing

- One important problem with linear probing is clustering — as collisions start to occur, then blocks of contiguous occupied bins (clusters) appear.
- And a quite unfortunate aspect is that the longer these clusters, the longer our searches or insertions (or deletions) will take (and remember that we wanted them to be constant time *and* fast!)

Hash Tables – Double hashing

- An even more unfortunate aspect is the fact that the longer these clusters, *the more likely* it will be that they will grow with each insertion!
  - This is because a new value inserted will make the cluster grow if the hash falls anywhere in the interval $[C_S-1, C_E+1]$, where C_S, C_E are the beginning and the end of the cluster, respectively.
    - Any hash that falls in the cluster will end up taking the position C_E+1, as a result of the linear probing.

Hash Tables – Double hashing

- One idea that could come to mind is to do linear probing using a jump size *p*; that is, if there is a collision, instead of skipping to the next bin to probe it, skip *p* bins forward and probe there.

Hash Tables – Double hashing

- One idea that could come to mind is to do linear probing using a jump size *p*; that is, if there is a collision, instead of skipping to the next bin to probe it, skip *p* bins forward and probe there.
- The bad news: It turns out that if the jump size is fixed, this does not make the slightest difference with respect to our "standard" linear probing (i.e., with jump size $p = 1$)

Hash Tables – Double hashing

- This is a direct consequence of the jump size being fixed.
  - Jump size different from one just makes it a bit more difficult to visualize, but the problem is exactly the same

Hash Tables – Double hashing

- So... What if we could choose a different jump size for each insertion?

Hash Tables – Double hashing

- So... What if we could choose a different jump size for each insertion?
- For example, the first insertion uses jump size 1, second insertion jump size 2, and so on...

Hash Tables – Double hashing

- So... What if we could choose a different jump size for each insertion?
- For example, the first insertion uses jump size 1, second insertion jump size 2, and so on...
  - Would this work, and avoid the issue of clustering?

Hash Tables – Double hashing

- So... What if we could choose a different jump size for each insertion?
- For example, the first insertion uses jump size 1, second insertion jump size 2, and so on...
  - Would this work, and avoid the issue of clustering?
  - We'll discuss in class why it doesn't!

Hash Tables – Double hashing

- However, if the jump size was a function of the value being inserted, things would work, right?

Hash Tables – Double hashing

- However, if the jump size was a function of the value being inserted, things would work, right?
- A function of the value being inserted ... sounds familiar, doesn't it?

Hash Tables – Double hashing

- However, if the jump size was a function of the value being inserted, things would work, right?
- A function of the value being inserted ... sounds familiar, doesn't it?
- It would be a bad idea to re-use the same hash function that we used to obtain the bin
  - However, we could use a second (different) hash function

Hash Tables – Double hashing

- We recall from two classes ago that we wanted to scramble the bits of the data and then select a subset of those bits (e.g., the *m* bits from the middle)
- What about taking advantage of the computation already done, and choose a *different* block of bits for the second hash function?

Hash Tables – Double hashing

- Let's look at an example, not with bits, but with something more human-brain-friendly:
  - The hash table uses size 10
  - For the hash function, multiply the value times 117 and keep the right-most digit
    - For the second hash function (jump size), just use the same result, and take the second digit

Hash Tables – Double hashing

- We'll insert values 14, 29, 43, 19, and 5 into the initially empty hash table:

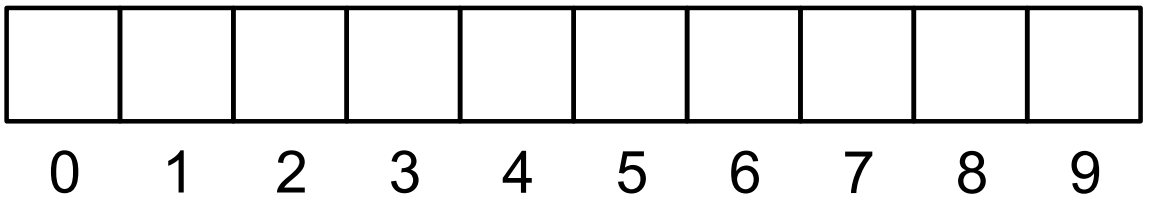

Hash Tables – Double hashing

- We'll insert values 14, 29, 43, 19, and 5 into the initially empty hash table:
  - 14×117 = 1638 ⇒ bin 8 (and jump size 3 — not relevant now, since this insertion causes no collision)

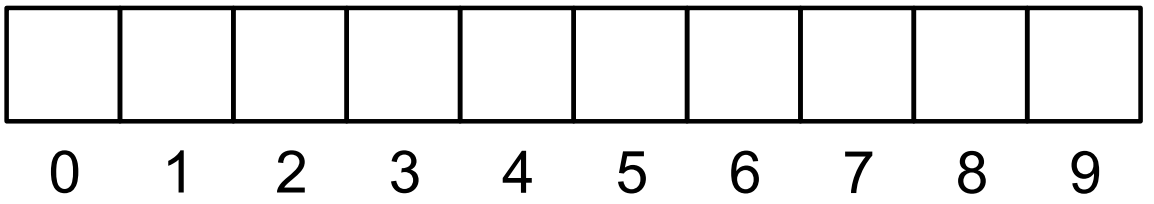

Hash Tables – Double hashing

- We'll insert values 14, 29, 43, 19, and 5 into the initially empty hash table:
  - 14×117 = 1638 ⇒ bin 8 (and jump size 3 — not relevant now, since this insertion causes no collision)

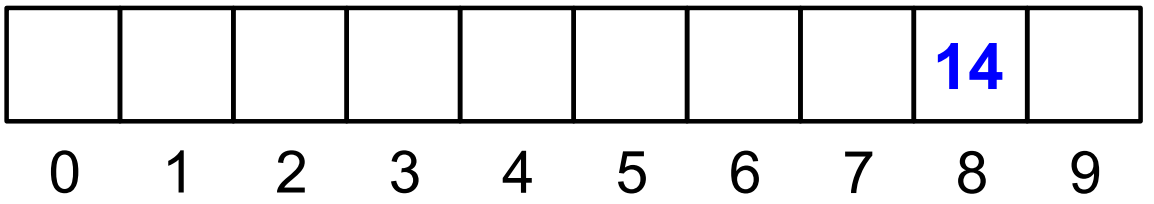

Hash Tables – Double hashing

- We'll insert values 14, 29, 43, 19, and 5 into the initially empty hash table:
  - 29×117 = 3393 ⇒ bin 3 (jump size not relevant)

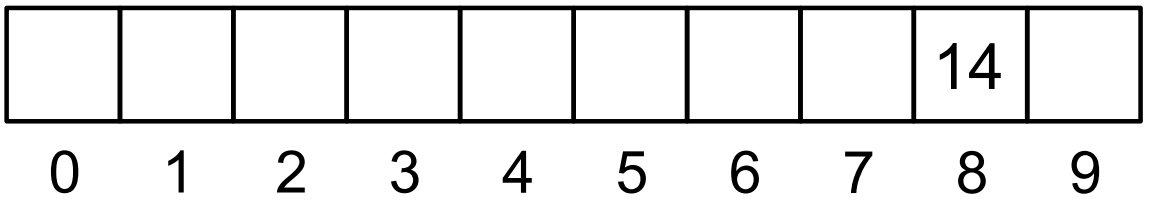

Hash Tables – Double hashing

- We'll insert values 14, 29, 43, 19, and 5 into the initially empty hash table:
  - 29×117 = 3393 ⇒ bin 3 (jump size not relevant)

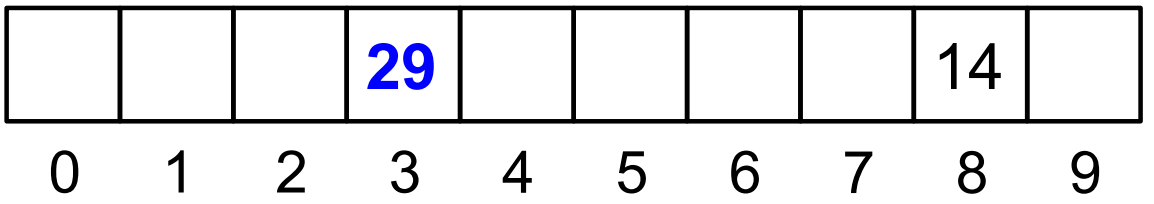

Hash Tables – Double hashing

- We'll insert values 14, 29, 43, 19, and 5 into the initially empty hash table:
  - 43×117 = 5031 ⇒ bin 1 (jump size not relevant)

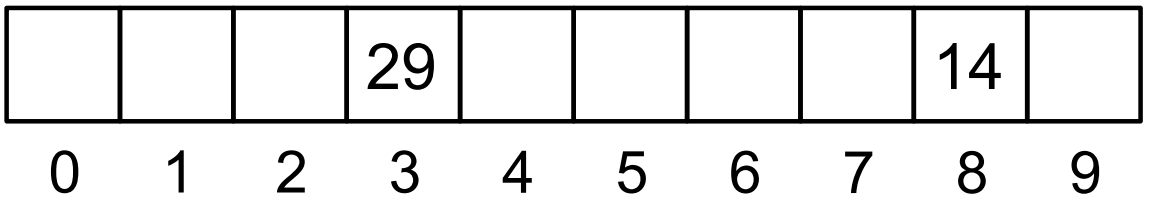

Hash Tables – Double hashing

- We'll insert values 14, 29, 43, 19, and 5 into the initially empty hash table:
  - 43×117 = 5031 ⇒ bin 1 (jump size not relevant)

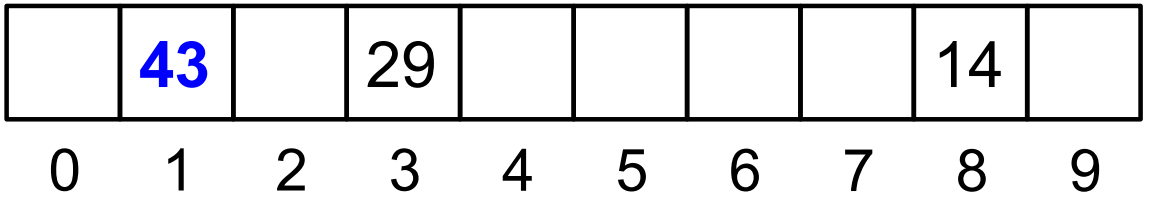

Hash Tables – Double hashing

- We'll insert values 14, 29, 43, 19, and 5 into the initially empty hash table:
  - 19×117 = 2223 ⇒ bin 3, causing a collision (jump size given by the second digit, 2)
    - Probe bin 3 + 2 — available, so we're done:

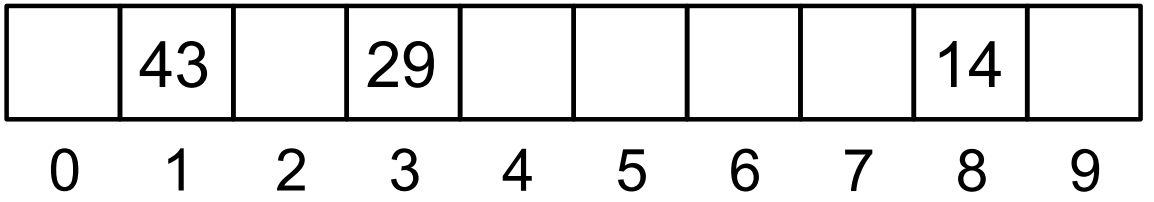

Hash Tables – Double hashing

- We'll insert values 14, 29, 43, 19, and 5 into the initially empty hash table:
  - 19×117 = 2223 ⇒ bin 3, causing a collision (jump size given by the second digit, 2)
    - Probe bin 3 + 2 — available, so we're done:

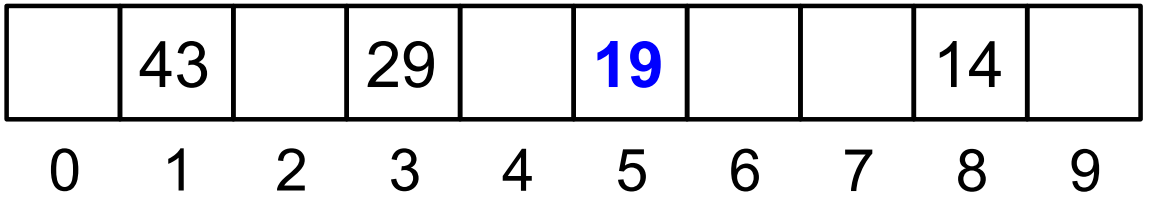

Hash Tables – Double hashing

- We'll insert values 14, 29, 43, 19, and 5 into the initially empty hash table:
  - 5×117 = 585 ⇒ bin 5, causing a collision (jump size given by the second digit, 8)
    - Where would this one end up?

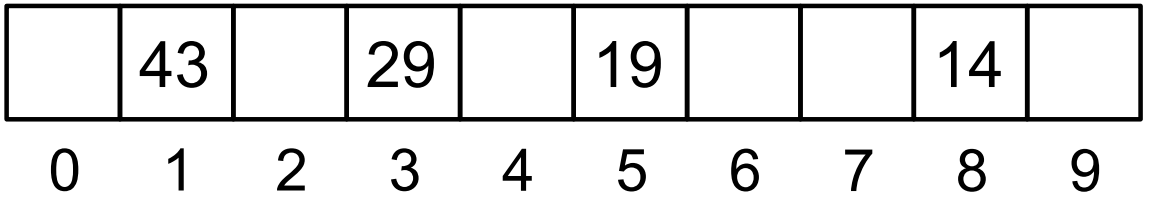

Hash Tables – Double hashing

- There's a big (read: BIG!) problem with this. Let's try inserting 59:

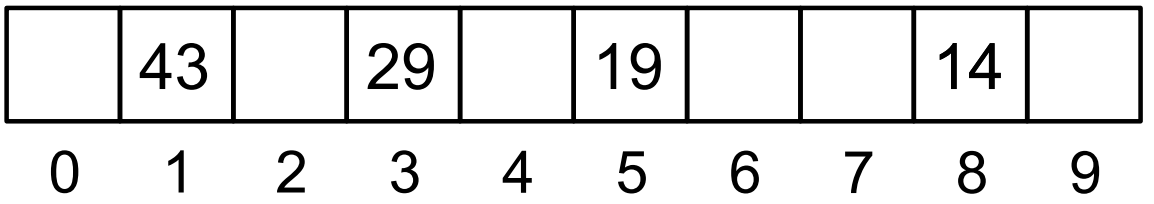

Hash Tables – Double hashing

- There's a big (read: BIG!) problem with this. Let's try inserting 59:
  - 59×117 = 6903 ⇒ bin 3, causing a collision, so we choose jump size ... Oops!

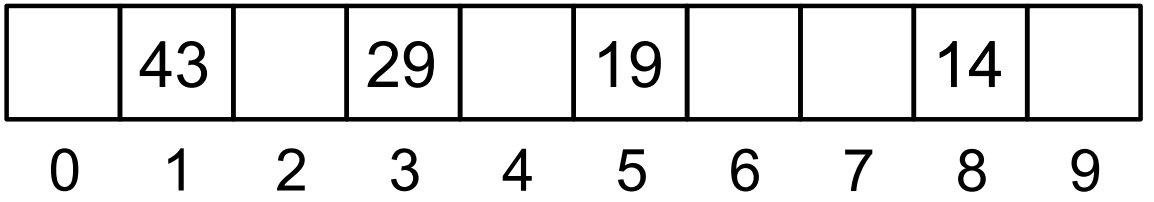

Hash Tables – Double hashing

- There's a big (read: BIG!) problem with this. Let's try inserting 59:
  - Ok, so we could fix this by not allowing the second hash function to take value 0 (*how do we do that?*)

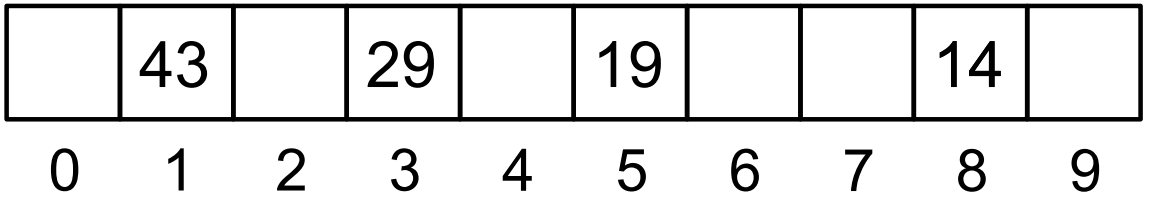

Hash Tables – Double hashing

- There's a big (read: BIG!) problem with this.
  - But that's not all — let's try inserting 74:

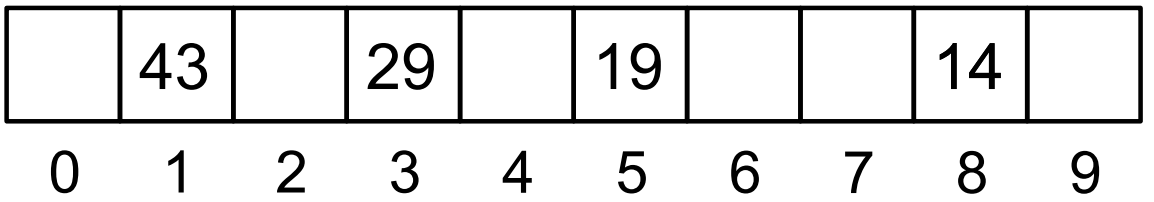

Hash Tables – Double hashing

- There's a big (read: BIG!) problem with this.
  - But that's not all — let's try inserting 74:
  - 74×117 = 8658 ⇒ bin 8, causing a collision, so we get jump size = 5, so we probe .... and... oops!

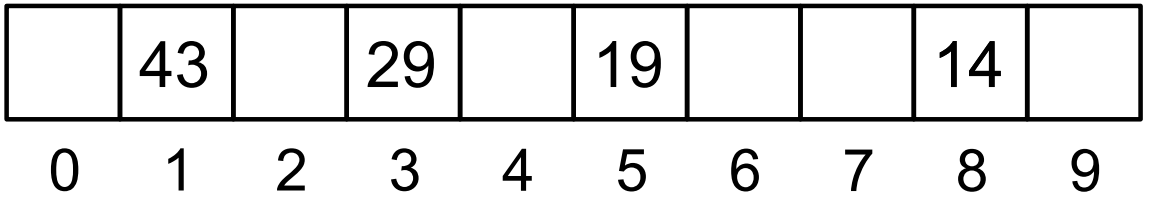

Hash Tables – Double hashing

- Why does this happens?
- How do we fix it?

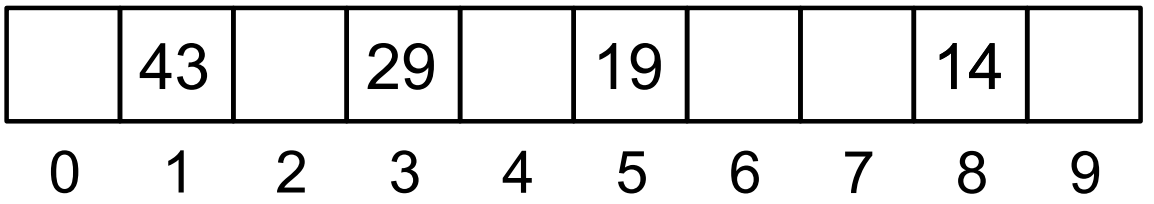

Hash Tables – Double hashing

- We'll discuss this in class.
- To be able to continue, let's say that if we choose sizes that are powers of two (and we always want to do that anyway), and we guarantee that the jump size is always an odd number, we avoid that issue.

Hash Tables – Double hashing

- How do we get a hash function that is guaranteed to be an odd value?

Hash Tables – Double hashing

- How do we get a hash function that is guaranteed to be an odd value?
  - Hint: what does an odd value look in binary?

Hash Tables – Double hashing

- Next, we'll look at removing values.

Hash Tables – Double hashing

- Can we use the same (efficient) trick as we did with linear probing?

Hash Tables – Double hashing

- Can we use the same (efficient) trick as we did with linear probing?
- How do we determine the sequence of elements to scan, if they're not consecutive?

Hash Tables – Double hashing

- Can we use the same (efficient) trick as we did with linear probing?
- How do we determine the sequence of elements to scan, if they're not consecutive?
- How do we check in constant time if a given element is going to be in trouble due to the hole left by the deletion?

Hash Tables – Double hashing

- Any ideas?

Hash Tables – Double hashing

- Any ideas?
  - Hint: we already saw the approach that we need here!

Summary

- During today's class, we discussed:
  - Clustering with linear probing
  - Double hashing:
    - Use one hash function to determine the bin
    - A second hash function determines the jump size for the probing sequence.
  - How to make the second hash suitable (typically, table size 2^m and jump size always odd)