CSE332: Data Abstractions

Lecture 10: More B Trees; Hashing

Dan Grossman

Spring 2012

B Tree Review

- M-ary tree with room for L data items at each leaf
- Order property:
  Subtree **between** keys *x* and *y* contains only data that is ≥ ***x*** and < ***y*** (notice the ≥)
- Balance property:
  All nodes and leaves at least half full, and all leaves at same height
- `find` and `insert` efficient
  - `insert` uses *splitting* to handle overflow, which may require splitting parent, and so on recursively

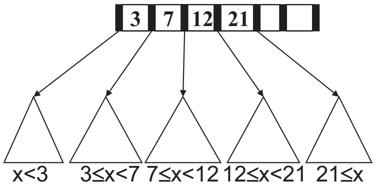

Can do a little better with insert

Eventually have to split up to the root (the tree will fill)

But can sometimes avoid splitting via *adoption*

- Change what leaf is correct by changing parent keys
- This idea "in reverse" is necessary in deletion (next)

Example:

Splitting

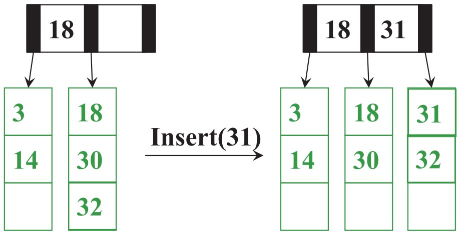

Adoption for insert

Eventually have to split up to the root (the tree will fill)

But can sometimes avoid splitting via *adoption*

- Change what leaf is correct by changing parent keys
- This idea "in reverse" is necessary in deletion (next)

Example:

Adoption

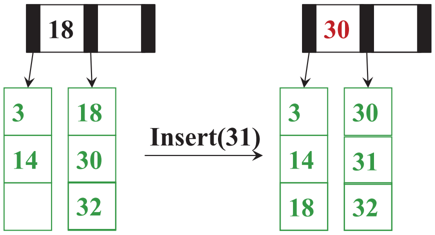

Deletion

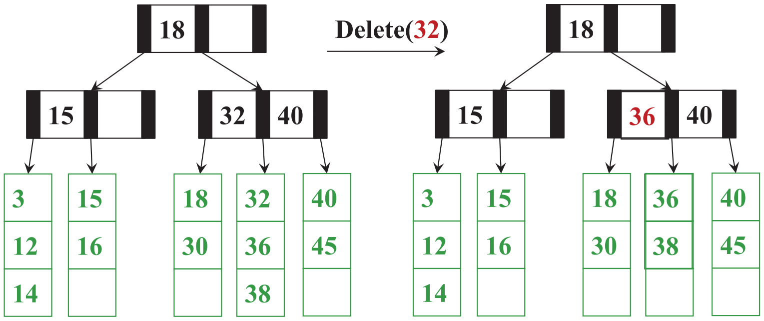

`M = 3 L = 3`

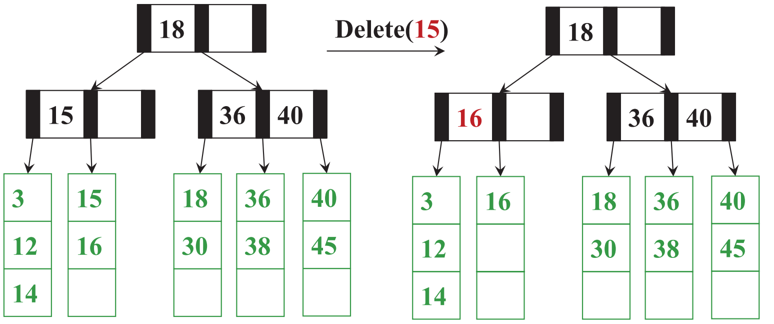

What's wrong?

Adopt from a neighbor!

`M = 3 L = 3`

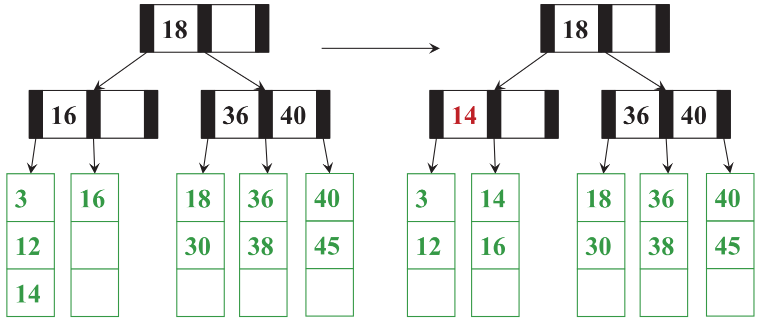

`M = 3 L = 3`

Delete(16)

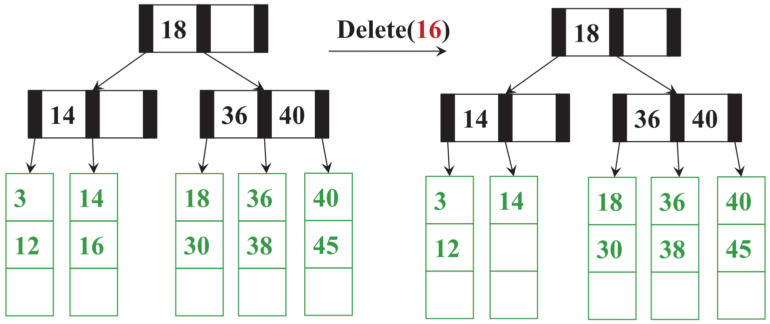

Uh-oh, neighbors at their minimum!

`M = 3 L = 3`

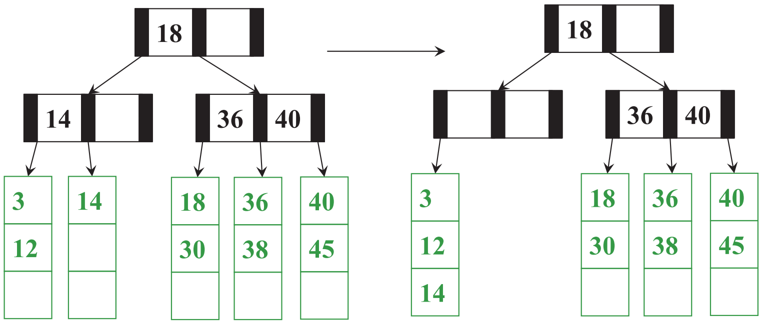

Move in together and remove leaf – now parent might underflow; it has neighbors

`M = 3 L = 3`

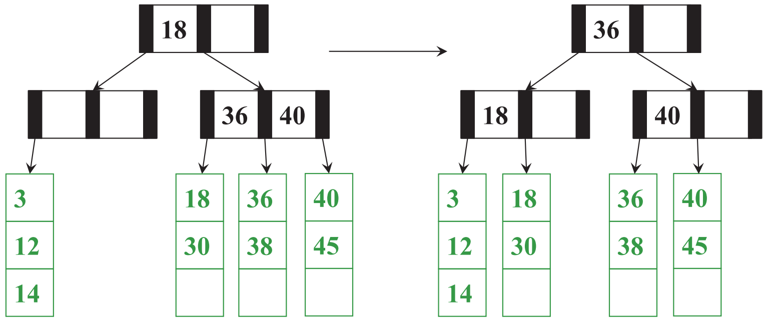

`M = 3 L = 3`

Delete(14)

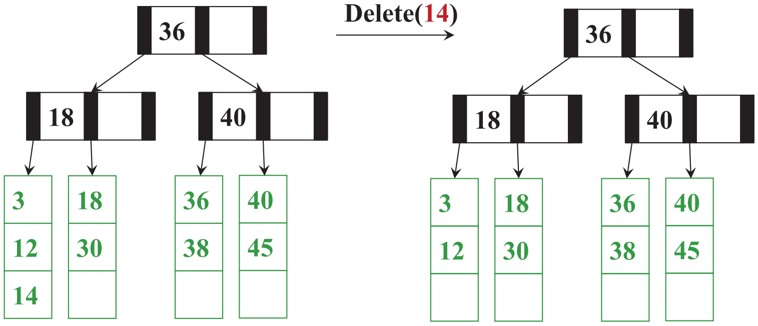

`M = 3 L = 3`

Delete(18)

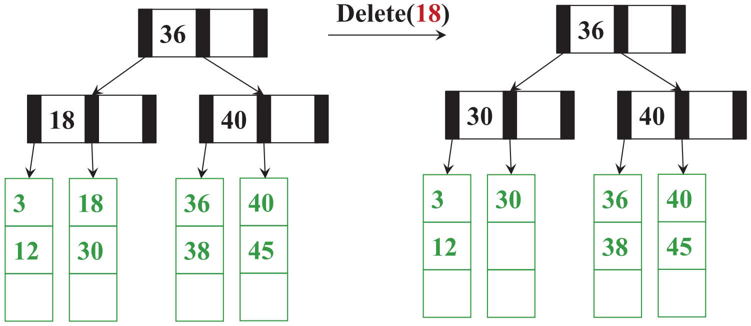

`M = 3 L = 3`

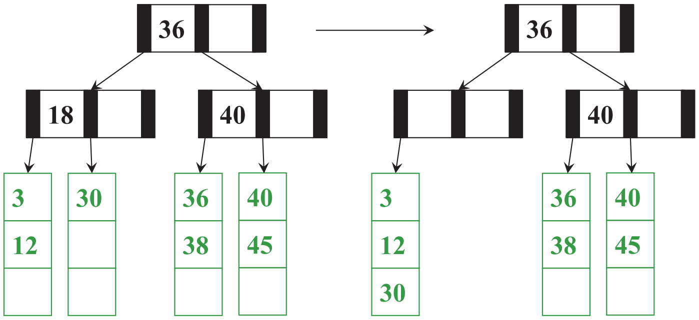

$M = 3$ $L = 3$

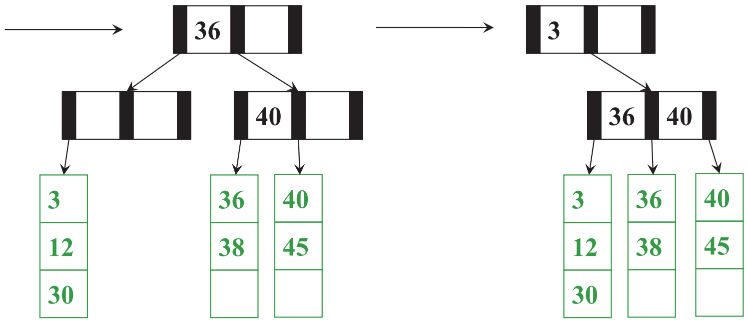

$M = 3$ $L = 3$

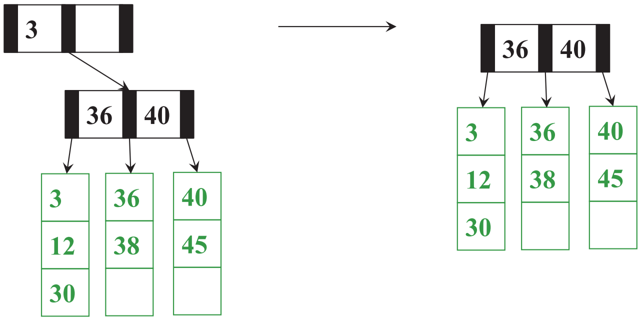

$M = 3$ $L = 3$

Deletion algorithm, part 1

1. Remove the data from its leaf

2. If the leaf now has $\lceil L/2 \rceil - 1$, *underflow!*
   - If a neighbor has > $\lceil L/2 \rceil$ items, *adopt* and update parent
   - Else *merge* node with neighbor
     - Guaranteed to have a legal number of items
     - Parent now has one less node

3. If step (2) caused the parent to have $\lceil M/2 \rceil - 1$ children, *underflow!*
   - …

Deletion algorithm continued

3. If an internal node has $\lceil M/2 \rceil - 1$ children
   - If a neighbor has > $\lceil M/2 \rceil$ items, *adopt* and update parent
   - Else *merge* node with neighbor
     - Guaranteed to have a legal number of items
     - Parent now has one less node, may need to continue underflowing up the tree

Fine if we merge all the way up through the root
- Unless the root went from 2 children to 1
- In that case, delete the root and make child the root
- This is the only case that decreases tree height

Worst-Case Efficiency of Delete

- Find correct leaf: $O(\log_2 M \log_M n)$
- Remove from leaf: $O(L)$
- Adopt from or merge with neighbor: $O(L)$
- Adopt or merge all the way up to root: $O(M \log_M n)$

Total: $O(L + M \log_M n)$

But it's not that bad:
- Merges are not that common
- Disk accesses are the name of the game: $O(\log_M n)$

B Trees in Java?

For most of our data structures, we have encouraged writing high-level, reusable code, such as in Java with generics

It is worthwhile to know enough about "how Java works" to understand why this is probably a bad idea for B trees
- If you just want a balanced tree with worst-case logarithmic operations, no problem
  - If *M*=3, this is called a 2-3 tree
  - If *M*=4, this is called a 2-3-4 tree
- Assuming our goal is efficient number of disk accesses
  - Java has many advantages, but it wasn't designed for this

The key issue is extra *levels of indirection*…

Naïve approach

Even if we assume data items have `int` keys, you cannot get the data representation you want for "really big data"

```
interface Keyed {
  int getKey();
}
class BTreeNode<E implements Keyed> {
  static final int M = 128;
  int[]          keys        = new int[M-1];
  BTreeNode<E>[] children    = new BTreeNode[M];
  int            numChildren = 0;
  …
}
class BTreeLeaf<E implements Keyed> {
  static final int L = 32;
  E[] data     = (E[])new Object[L];
  int numItems = 0;
  …
}
```

What that looks like

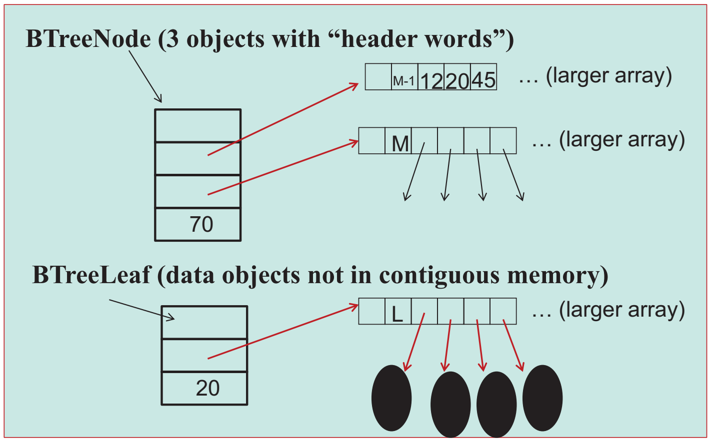

BTreeNode (3 objects with "header words")

BTreeLeaf (data objects not in contiguous memory)

The moral

- The point of B trees is to keep related data in contiguous memory
- All the red references on the previous slide are inappropriate
  - As minor point, beware the extra "header words"
- But that's "the best you can do" in Java
  - Again, the advantage is generic, reusable code
  - But for your performance-critical web-index, not the way to implement your B-Tree for terabytes of data
- Other languages (e.g., C++) have better support for "flattening objects into arrays"
- Levels of indirection matter!

Conclusion: Balanced Trees

- *Balanced* trees make good dictionaries because they guarantee logarithmic-time **`find`**, **`insert`**, and **`delete`**
  - Essential and beautiful computer science
  - But only if you can maintain balance within the time bound
- AVL trees maintain balance by tracking height and allowing all children to differ in height by at most 1
- B trees maintain balance by keeping nodes at least half full and all leaves at same height
- Other great balanced trees (see text; worth knowing they exist)
  - Red-black trees: all leaves have depth within a factor of 2
  - Splay trees: self-adjusting; amortized guarantee; no extra space for height information

Motivating Hash Tables

For dictionary with *n* key/value pairs

	`insert`	**`find`**	**`delete`**
• Unsorted linked-list	$O(1)$	$O(n)$	$O(n)$
• Unsorted array	$O(1)$	$O(n)$	$O(n)$
• Sorted linked list	$O(n)$	$O(n)$	$O(n)$
• Sorted array	$O(n)$	$O(\log n)$	$O(n)$
• *Balanced* tree	$O(\log n)$	$O(\log n)$	$O(\log n)$
• Magic array	$O(1)$	$O(1)$	$O(1)$

Sufficient "magic":
- Compute array index for an item in $O(1)$ time [doable]
- Have a different index for every item [magic]

Hash Tables

- Aim for constant-time (i.e., *O*(1)) **`find`**, **`insert`**, and **`delete`**
  - "On average" under some often-reasonable assumptions
- A hash table is an array of some fixed size
- Basic idea:

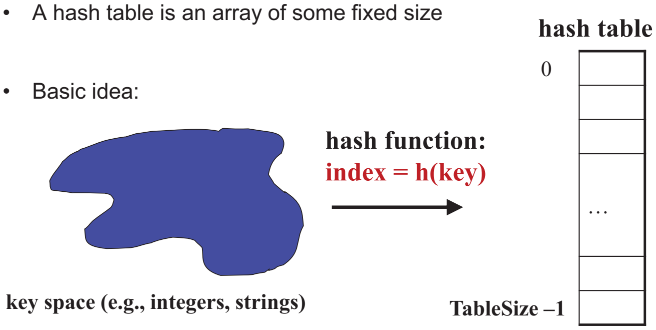
hash table
0
…
TableSize –1
hash function:
index = h(key)
key space (e.g., integers, strings)

Hash Tables vs. Balanced Trees

- In terms of a Dictionary ADT for just **`insert`**, **`find`**, **`delete`**, hash tables and balanced trees are just different data structures
  - Hash tables *O*(1) on average (*assuming* few collisions)
  - Balanced trees *O*(`log` *n*) worst-case
- Constant-time is better, right?
  - Yes, but you need "hashing to behave" (must avoid collisions)
  - Yes, but **`findMin`**, **`findMax`**, **`predecessor`**, and **`successor`** go from *O*(`log` *n*) to *O*(*n*), **`printSorted`** from *O*(*n*) to *O*(*n* `log` *n*)
    - Why your textbook considers this to be a different ADT
    - Not so important to argue over the definitions

Hash Tables

- There are *m* possible keys (*m* typically large, even infinite)
- We expect our table to have only *n* items
- *n* is much less than *m* (often written *n* << *m*)

Many dictionaries have this property

- Compiler: All possible identifiers allowed by the language vs. those used in some file of one program
- Database: All possible student names vs. students enrolled
- AI: All possible chess-board configurations vs. those considered by the current player
- …

Hash functions

An ideal hash function:

- Is fast to compute
- "Rarely" hashes two "used" keys to the same index
  - Often impossible in theory; easy in practice
  - Will handle *collisions* in next lecture

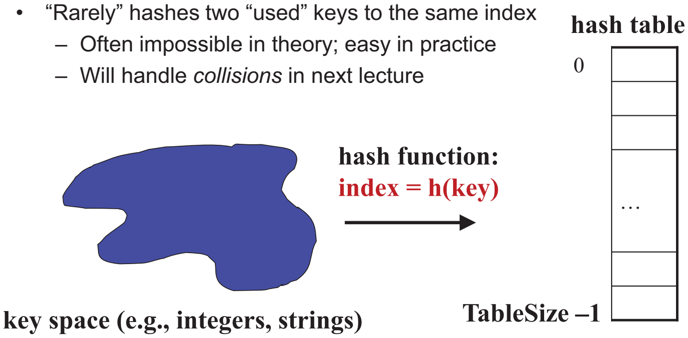
hash table
0
…
TableSize –1
hash function:
index = h(key)
key space (e.g., integers, strings)

Who hashes what?

- Hash tables can be generic
  - To store elements of type **`E`**, we just need **`E`** to be:
    1. Comparable: order any two **`E`** (as with all dictionaries)
    2. Hashable: convert any **`E`** to an **`int`**
- When hash tables are a reusable library, the division of responsibility generally breaks down into two roles:

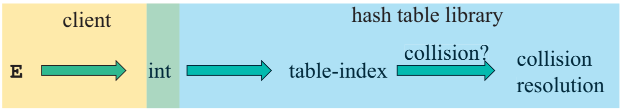
client
hash table library
`E`
int
table-index
collision?
collision resolution

- We will learn both roles, but most programmers "in the real world" spend more time as clients while understanding the library

More on roles

Some ambiguity in terminology on which parts are "hashing"

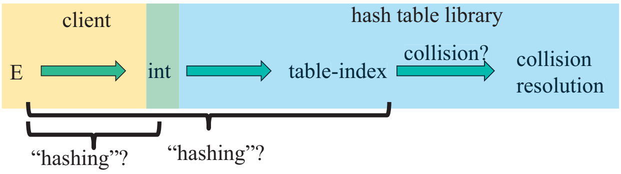
client
hash table library
E
int
table-index
collision?
collision resolution
"hashing"?
"hashing"?

Two roles must both contribute to minimizing collisions (heuristically)

- Client should aim for different ints for expected items
  - Avoid "wasting" any part of **`E`** or the 32 bits of the **`int`**
- Library should aim for putting "similar" **`int`**s in different indices
  - Conversion to index is almost always "mod table-size"
  - Using prime numbers for table-size is common

What to hash?

We will focus on the two most common things to hash: ints and strings

- If you have objects with several fields, it is usually best to have most of the “identifying fields” contribute to the hash to avoid collisions
- Example:

```
class Person {
   String first; String middle; String last;
   Date birthdate;
}
```

- An inherent trade-off: hashing-time vs. collision-avoidance
  - Bad idea(?): Only use first name
  - Good idea(?): Only use middle initial
  - Admittedly, what-to-hash is often unprincipled ☹

Hashing integers

- key space = integers
- Simple hash function:
  `h(key) = key % TableSize`
  - Client: **`f(x) = x`**
  - Library **`g(x) = x % TableSize`**
  - Fairly fast and natural
- Example:
  - **`TableSize`** = 10
  - Insert 7, 18, 41, 34, 10
  - (As usual, ignoring data “along for the ride”)

0	
1	
2	
3	
4	
5	
6	
7	
8	
9	

Hashing integers

- key space = integers
- Simple hash function:
  `h(key) = key % TableSize`
  - Client: **`f(x) = x`**
  - Library **`g(x) = x % TableSize`**
  - Fairly fast and natural
- Example:
  - **`TableSize`** = 10
  - Insert 7, 18, 41, 34, 10
  - (As usual, ignoring data “along for the ride”)

0	10
1	41
2	
3	
4	34
5	
6	
7	7
8	18
9	

Collision-avoidance

- With “**`x % TableSize`**” the number of collisions depends on
  - the ints inserted (obviously)
  - **`TableSize`**
- Larger table-size tends to help, but not always
  - Example: 70, 24, 56, 43, 10
    with **`TableSize`** = 10 and **`TableSize`** = 60
- Technique: Pick table size to be prime. Why?
  - Real-life data tends to have a pattern
  - “Multiples of 61” are probably less likely than “multiples of 60”
  - Next time we’ll see that one collision-handling strategy does *provably* well with prime table size

---

More on prime table size

If **`TableSize`** is 60 and…

- Lots of data items are multiples of 5, wasting 80% of table
- Lots of data items are multiples of 10, wasting 90% of table
- Lots of data items are multiples of 2, wasting 50% of table

If **`TableSize`** is 61…

- Collisions can still happen, but 5, 10, 15, 20, … will fill table
- Collisions can still happen but 10, 20, 30, 40, … will fill table
- Collisions can still happen but 2, 4, 6, 8, … will fill table

In general, if **`x`** and **`y`** are “co-prime” (means **`gcd(x,y)==1`**), then
`(a * x) % y == (b * x) % y` if and only if **`a % y == b % y`**

- So good to have a **`TableSize`** that has no common factors with any “likely pattern” **`x`**

Okay, back to the client

- If keys aren’t **`ints`**, the client must convert to an **`int`**
  - Trade-off: speed and distinct keys hashing to distinct **`ints`**
- Very important example: Strings
  - Key space K $= s_0 s_1 s_2 \ldots s_{m-1}$
    - (where s_i are chars: $s_i \in [0,52]$ or $s_i \in [0,256]$ or $s_i \in [0,2^{16}]$)
  - Some choices: Which avoid collisions best?

1. h(K) = s_0 % TableSize
2. h(K) = $\left(\sum_{i=0}^{m-1} s_i\right)$ % TableSize
3. h(K) = $\left(\sum_{i=0}^{k-1} s_i \cdot 37^i\right)$ % TableSize

Specializing hash functions

How might you hash differently if all your strings were web addresses (URLs)?

Combining hash functions

A few rules of thumb / tricks:

1. Use all 32 bits (careful, that includes negative numbers)
2. Use different overlapping bits for different parts of the hash
   - This is why a factor of 37^i works better than 256^i
   - Example: “abcde” and “ebcda”
3. When smashing two hashes into one hash, use bitwise-xor
   - bitwise-and produces too many 0 bits
   - bitwise-or produces too many 1 bits
4. Rely on expertise of others; consult books and other resources
5. If keys are known ahead of time, choose a *perfect hash*

One expert suggestion

- int result = 17;
- foreach field f
  - int fieldHashcode =
    - boolean: (f ? 1: 0)
    - byte, char, short, int: (int) f
    - long: (int) (f ^ (f >>> 32))
    - float: Float.floatToIntBits(f)
    - double: Double.doubleToLongBits(f), then above
    - Object: object.hashCode( )
  - result = 31 * result + fieldHashcode