CS 240 – Data Structures and Data Management

Module 4: Dictionaries

Armin Jamshidpey, Éric Schost Based on lecture notes by many previous cs240 instructors

David R. Cheriton School of Computer Science, University of Waterloo

Spring 2025

version 2025-06-01 21:23

Outline

Dictionaries and Balanced Search Trees

- ADT Dictionary
- Binary Search Trees
- AVL Trees
- Insertion in AVL Trees
- Restructuring a BST: Rotations
- AVL insertion revisited
- Deletion in AVL Trees

Outline

4 Dictionaries and Balanced Search Trees

- ADT Dictionary
- Binary Search Trees
- AVL Trees
- Insertion in AVL Trees
- Restructuring a BST: Rotations
- AVL insertion revisited
- Deletion in AVL Trees

Review: ADT Dictionary

Dictionary: A collection of items, each of which contains

- a *key*
- some *data* (the "value")

and is called a *key-value pair* (KVP). Keys can be compared and are (typically) unique.

Operations:

- search(k) (also called lookup(k))
- insert(k, v)
- delete(k) (also called remove(k))
- optional: successor, merge, is-empty, size, etc.

Examples: symbol table, license plate database

Review: Elementary realizations

Common assumptions:

- Dictionary has *n* KVPs
- Each KVP uses constant space (if not, the "value" could be a pointer)
- Keys can be compared in constant time

Review: Elementary realizations

Common assumptions:

- Dictionary has n KVPs
- Each KVP uses constant space (if not, the "value" could be a pointer)
- Keys can be compared in constant time

We commonly make one more assumption (to keep pseudo-code simple):

• Dictionary is non-empty both before and after operation.

(In a real-life implementation you would have to treat these special cases.)

Review: Elementary realizations

Common assumptions:

- Dictionary has n KVPs
- Each KVP uses constant space (if not, the "value" could be a pointer)
- Keys can be compared in constant time

We commonly make one more assumption (to keep pseudo-code simple):

• Dictionary is non-empty both before and after operation.

(In a real-life implementation you would have to treat these special cases.)

Easy realizations:

	search	insert	delete
unsorted list/array	$\Theta(n)$	$\Theta(1)$	$\Theta(1)$
sorted array	$\Theta(\log n)$	$\Theta(n)$	$\Theta(n)$
binary search tree	$\Theta(height)$	$\Theta(height)$	$\Theta(height)$

Review: Binary search



Only applies to a *sorted array*:



We will return to binary search (and sometimes improve it!) later.

Outline

Dictionaries and Balanced Search Trees

ADT Dictionary

• Binary Search Trees

- AVL Trees
- Insertion in AVL Trees
- Restructuring a BST: Rotations
- AVL insertion revisited
- Deletion in AVL Trees

Review: Binary search trees

Structure Binary tree: all nodes have two (possibly empty) subtrees Every node stores a KVP Empty subtrees usually not shown

Ordering Every key k in *T.left* is less than the root key. Every key k in *T.right* is greater than the root key.



(In our examples we only show the keys, and we show them directly in the node. A more accurate picture would be $(\bullet, \bullet, \bullet, \bullet, \bullet)$ (key = 15, <other info>)

CS240 - Module 4

BST::search(k) Start at root, compare k to current node's key. Stop if found or subtree is empty, else recurse at subtree.



BST::search(k) Start at root, compare k to current node's key. Stop if found or subtree is empty, else recurse at subtree.



BST::search(k) Start at root, compare k to current node's key. Stop if found or subtree is empty, else recurse at subtree.



BST::search(k) Start at root, compare k to current node's key. Stop if found or subtree is empty, else recurse at subtree.



BST::search(k) Start at root, compare k to current node's key. Stop if found or subtree is empty, else recurse at subtree. BST::insert(k, v) Search for k, then insert (k, v) as new node

Example: *BST::insert*(24, v)



- First search for the node x that contains the key.
- If x is a **leaf** (both subtrees are empty), delete it.



- First search for the node x that contains the key.
- If x is a **leaf** (both subtrees are empty), delete it.



- First search for the node x that contains the key.
- If x is a **leaf** (both subtrees are empty), delete it.
- If x has one non-empty subtree, move child up



- First search for the node x that contains the key.
- If x is a **leaf** (both subtrees are empty), delete it.
- If x has one non-empty subtree, move child up



- First search for the node x that contains the key.
- If x is a **leaf** (both subtrees are empty), delete it.
- If x has one non-empty subtree, move child up
- Else, swap key at x with key at **successor** node and then delete that node.

(Successor: next-smallest among all keys in the dictionary.)



- First search for the node x that contains the key.
- If x is a **leaf** (both subtrees are empty), delete it.
- If x has one non-empty subtree, move child up
- Else, swap key at x with key at **successor** node and then delete that node.

(Successor: next-smallest among all keys in the dictionary.)



- First search for the node x that contains the key.
- If x is a **leaf** (both subtrees are empty), delete it.
- If x has one non-empty subtree, move child up
- Else, swap key at x with key at **successor** node and then delete that node.

(Successor: next-smallest among all keys in the dictionary.)



Height of a BST

BST::search, BST::insert, BST::delete all have cost $\Theta(h)$, where h = maximum number for which level h contains nodes. Single-node tree has height 0, empty tree has height -1

If n items are inserted one-at-a-time, how big is h?

• Worst-case:
$$n-1=\Theta(n)$$

 Best-case: Θ(log n). Any binary tree with n nodes has height h ≥ log(n + 1) − 1 (See Module 2).

Height of a BST

BST::search, BST::insert, BST::delete all have cost $\Theta(h)$, where h = maximum number for which level h contains nodes. Single-node tree has height 0, empty tree has height -1

If n items are inserted one-at-a-time, how big is h?

• Worst-case:
$$n-1=\Theta(n)$$

 Best-case: Θ(log n). Any binary tree with n nodes has height h ≥ log(n + 1) − 1 (See Module 2).

Goal: Create subclasses of BSTs where the height is *always* good.

- Impose a structural property.
- Argue that the property implies logarithmic height.
- Discuss how to maintain the property during operatons.

Outline

Dictionaries and Balanced Search Trees

- ADT Dictionary
- Binary Search Trees

AVL Trees

- Insertion in AVL Trees
- Restructuring a BST: Rotations
- AVL insertion revisited
- Deletion in AVL Trees

AVL trees

Introduced by Adel'son-Vel'skiĭ and Landis in 1962, an **AVL tree** is a BST with an additional **height-balance property** at every node:

The heights of the left and right subtree differ by at most 1.

Rephrase: If node z has left subtree L and right subtree R, then

height(R) - height(L) must be in $\{-1, 0, 1\}$

AVL trees

Introduced by Adel'son-Vel'skiĭ and Landis in 1962, an **AVL tree** is a BST with an additional **height-balance property** at every node:

The heights of the left and right subtree differ by at most 1.

Rephrase: If node z has left subtree L and right subtree R, then

height(R) - height(L) must be in $\{-1, 0, 1\}$

Things to show:

- This structural condition implies logarithmic height.
- After *insert* and *delete*, we can restore the structural condition within logarithmic time.
 - ► For this, we need to store at each node *z* the height of the subtree rooted at it.

AVL tree example

(The lower numbers indicate the height of the subtree.)



AVL tree example

Alternative: store height-difference (instead of height) at each node.



- Saves space (2 bits vs. 1 integer per node)
- $\bullet\,$ Pseudo-code gets a lot more complicated \rightsquigarrow not done here

A. Jamshidpey, É. Schost (CS-UW)

CS240 - Module 4

Spring 2025

AVL trees: Height

Theorem: The height of an AVL tree on *n* nodes is in $\Theta(\log n)$ \Rightarrow search, BST::insert, BST::delete all cost $\Theta(\log n)$ in the worst case!

Proof:

- Define N(h) to be the *least* number of nodes in a height-*h* AVL tree.
- What is a recurrence relation for N(h)?
- What does this recurrence relation resolve to?

Outline

Dictionaries and Balanced Search Trees

- ADT Dictionary
- Binary Search Trees
- AVL Trees

Insertion in AVL Trees

- Restructuring a BST: Rotations
- AVL insertion revisited
- Deletion in AVL Trees

AVL trees: Insertion

- To perform AVL::insert(k, v):
 - First, insert (k, v) with the usual BST insertion.
 - We assume that this returns the new leaf z where the key was stored.
 - Then, move up the tree from z.

(We assume for this that we have parent-links. This can be avoided if BST::insert returns the full path to z.

• Update height (easy to do in constant time):

- set-height-from-subtrees(u) 1. u.height $\leftarrow 1 + \max\{u.left.height, u.right.height\}$
- If the height difference becomes ± 2 at node z, then z is unbalanced. Must re-structure the tree to rebalance.








Outline

Dictionaries and Balanced Search Trees

- ADT Dictionary
- Binary Search Trees
- AVL Trees
- Insertion in AVL Trees

• Restructuring a BST: Rotations

- AVL insertion revisited
- Deletion in AVL Trees

Changing structure without changing order

Note: There are many different BSTs with the same keys.



Goal: Change the *structure* locally without changing the *order*.

Longterm goal: Restructure such that the subtree becomes balanced.

Right Rotation

This is a **right rotation** on node *z*:



Right Rotation

This is a **right rotation** on node *z*:



Note: Only O(1) links are changed. Useful to fix left-left imbalance.

Why do we call this a rotation?



Why do we call this a rotation?



Why do we call this a rotation?



Right Rotation: Pseudocode

```
rotate-right(z)
1. c \leftarrow z left
2. // fix links connecting to above
3. c.parent \leftarrow (p \leftarrow z.parent)
4. if p = NULL then root \leftarrow c else
5.
          if p.left = z then p.left \leftarrow c else p.right \leftarrow c
6. // actual rotation
7. z.left \leftarrow c.right, c.right.parent \leftarrow z
8. c.right \leftarrow z, z.parent \leftarrow c
9. set-height-from-subtrees(z), set-height-from-subtrees(c)
10. return c // returns new root of subtree
```

Run-time: O(1)

Left Rotation

Symmetrically, this is a **left rotation** on node *z*:



Again, only O(1) links need to be changed. Useful to fix right-right imbalance.

Double Right Rotation

This is a **double right rotation** on node *z*:



First, a left rotation at c.

Double Right Rotation

This is a **double right rotation** on node *z*:



First, a left rotation at c. Second, a right rotation at z.

Double Left Rotation

Symmetrically, there is a **double left rotation** on node *z*:



First, a right rotation at c. Second, a left rotation at z.

Outline

Dictionaries and Balanced Search Trees

- ADT Dictionary
- Binary Search Trees
- AVL Trees
- Insertion in AVL Trees
- Restructuring a BST: Rotations
- AVL insertion revisited
- Deletion in AVL Trees

AVL tree insertion: Example revisited

Example: *AVL::insert*(8)



AVL tree insertion revisited

- Imbalance at z: do (single or double) rotation
- Choose c as child where subtree has bigger height.

```
AVL::insert(k, v)
1. z \leftarrow BST::insert(k, v)
                                       // new leaf with k
    while (z is not NULL)
2.
3.
          if (|z.left.height - z.right.height| > 1) then
               Let c be taller child of z
4
               Let g be taller child of c (so grandchild of z)
5.
               z \leftarrow restructure(g, c, z) // \text{ see later}
6
7.
               break
                              // can argue that we are done
8.
        set-height-from-subtrees(z)
9.
          z \leftarrow z.parent
```

Can argue: For insertion *one* rotation restores all heights of subtrees. \Rightarrow No further imbalances, can stop checking.

Fixing a slightly-unbalanced AVL tree



Rule: The middle key of g, c, z becomes the new root.

AVL tree insertion: Example revisited



AVL tree insertion: Example revisited

















Outline

Dictionaries and Balanced Search Trees

- ADT Dictionary
- Binary Search Trees
- AVL Trees
- Insertion in AVL Trees
- Restructuring a BST: Rotations
- AVL insertion revisited
- Deletion in AVL Trees

AVL trees: Deletion

Remove the key k with BST::delete.

Find node where *structural* change happened.

(This is not necessarily near the node that had k.) Go back up to root, update heights, and rotate if needed.

```
AVL::delete(k)
1. z \leftarrow BST::delete(k)
2. // Assume z is the parent of the BST node that was removed
    while (z is not NULL)
3.
         if (|z.left.height - z.right.height| > 1) then
4.
              Let c be taller child of z
5.
6.
              Let g be taller child of c (break ties to avoid double rotation)
7.
              z \leftarrow restructure(g, c, z)
8. // Always continue up the path
9.
        set-height-from-subtrees(z)
10.
         z \leftarrow z.parent
```











A single *restructure* is not enough to restore all balances.

CS240 - Module 4





Important: Ties *must* be broken to avoid double rotation. Consider again the above example. If we applied double-rotation:



Important: Ties *must* be broken to avoid double rotation. Consider again the above example. If we applied double-rotation:



Violation is *below* where we check further.

AVL trees: Summary

search: Just like in BSTs, costs $\Theta(height)$

insert: BST::insert, then check & update along path to new leaf

- total cost $\Theta(height)$
- restructure will be called at most once.

delete: BST::delete, then check & update along path to deleted node

- total cost $\Theta(height)$
- restructure may be called $\Theta(height)$ times.

Worst-case cost for all operations is $\Theta(height) = \Theta(\log n)$.

- In practice, the constant is quite large.
- Other realizations of ADT Dictionary are better in practice (\rightarrow later)