

## Chapter 4 Trees

1

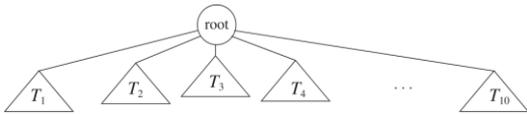
### Introduction

- for large input, even linear access time may be prohibitive
- we need data structures that exhibit average running times closer to  $O(\log N)$
- binary search tree

2

### Terminology

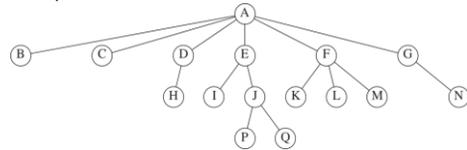
- recursive definition of **tree**
  - collection of nodes (may be empty)
  - distinguished node,  $r$ , is the **root**
  - zero or more nonempty subtrees  $T_1, T_2, \dots, T_k$ , each of whose roots are connected by a directed **edge** from  $r$
- root of each subtree is a **child** of  $r$
- $r$  is the **parent** of each subtree
- tree of  $N$  nodes has  $N - 1$  edges



3

### Terminology

- example tree

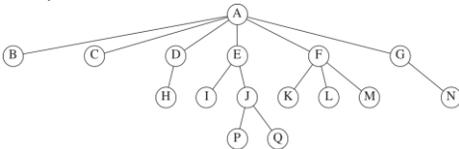


- nodes with no children are called **leaves** (e.g., B, C, H, I, P, Q, K, L, M, N)
- nodes with the same parent are **siblings** (e.g., K, L, M)
- parent, grandparent, grandchild, ancestor, descendant, proper ancestor, proper descendant**

4

### Terminology

- example tree

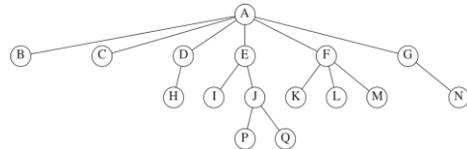


- path** from  $n_i$  to  $n_k$  is a sequence of nodes  $n_1, n_2, \dots, n_k$  where  $n_1$  is the parent of  $n_{i+1}$  for  $1 \leq i < k$
- length** of path is number of edges on path ( $k - 1$ )
  - path of length 0 from every node to itself
  - exactly one path from the root to each node

5

### Terminology

- example tree



- depth** from  $n_i$  is the length of the unique path from the root to  $n_i$ 
  - root is at depth 0
- height** of  $n_i$  is the length of the longest path from  $n_i$  to a leaf
  - all leaves at height 0
  - height of the tree is equal to the height of the root

6

### Terminology

-example tree

```

graph TD
    A((A)) --- B((B))
    A --- C((C))
    A --- D((D))
    A --- E((E))
    A --- F((F))
    A --- G((G))
    E --- H((H))
    E --- I((I))
    E --- J((J))
    E --- K((K))
    E --- L((L))
    E --- M((M))
    J --- P((P))
    J --- Q((Q))
    G --- N((N))
    
```

-E is at depth 1 and height 2  
 -E is at depth 1 and height 1  
 -depth of tree is 3

7

### Implementation of Trees

-each node could have data and a link to each child  
 -number of children is unknown and may be large, which could lead to wasted space  
 -instead, keep children in a linked list

```

1 struct TreeNode
2 {
3     Object element;
4     TreeNode *firstChild;
5     TreeNode *nextSibling;
6 };
    
```

-null links not shown

8

### Tree Traversals with Application

-many applications for trees  
 -subdirectory structure in Unix  
 -pathname built into tree

9

### Tree Traversals with Application

-goal: list all files in a directory  
 -depth denoted by tabs  
 -begins at root

```

void FileSystem::listAll( int depth = 0 ) const
{
1     printName( depth ); // Print the name of the object
2     if( isDirectory( ) )
3         for each file c in this directory ( for each child )
4         c.listAll( depth + 1 );
}
    
```

10

### Tree Traversals with Application

-code prints directories/files in preorder traversal  
 -runs in  $O(N)$

11

### Tree Traversals with Application

-for postorder traversal, numbers in parentheses represent the number of disk blocks for each file

12

## Tree Traversals with Application

- size method to find number of blocks for each file
- directories use 1 block of space

```

int FileSystem::size() const
{
    int totalSize = sizeOfThisFile( );
    if( !isDirectory( ) )
        for each file c in this directory (for each child)
            totalSize += c.size( );
    return totalSize;
}

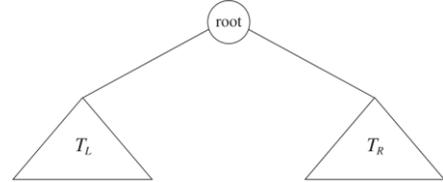
```

chl.r	3
chl.r	2
chl.r	4
book	10
syt.r	1
fall	2
syt.r	5
spr	6
syt.r	2
sum	3
cap3030	12
course	13
junk	6
mark	30
junk	8
alex	9
work	1
grades	3
pragl.r	4
pragl.r	1
fall	9
pragl.r	2
pragl.r	7
grades	9
fall	19
cap3212	19
course	30
bill	32
/sur	72

13

## Binary Trees

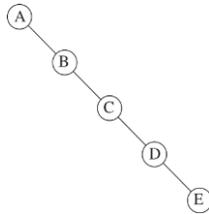
- in binary trees, nodes can have no more than two children
- binary tree below consists of a root and two subtrees,  $T_L$  and  $T_R$ , both of which could possibly be empty



14

## Binary Trees

- depth of a binary tree is considerably smaller than  $\underline{N}$
- average depth is  $O(\sqrt{N})$
- average depth for a binary search tree is  $O(\log N)$
- depth can be as large as  $\underline{N-1}$



15

## Binary Tree Implementation

- since a binary tree has two children at most, we can keep direct links to each of them
- element plus two pointers, left and right

```

struct BinaryNode
{
    Object element; // The data in the node
    BinaryNode *left; // Left child
    BinaryNode *right; // Right child
};

```

- drawn with circles and lines (graph)
- many applications, including compiler design

16

## Tree Traversals

- easy to list all elements of a binary search tree in sorted order
- inorder traversal
- postorder traversal
- preorder traversal
- implemented with recursive functions
- all  $O(N)$

17

## Tree Traversals

- inorder traversal

```

1 /**
2  * Print the tree contents in sorted order.
3  */
4 void printTree( ostream & out = cout ) const
5 {
6     if( isEmpty( ) )
7         out << "Empty tree" << endl;
8     else
9         printTree( root, out );
10 }
11
12 /**
13  * Internal method to print a subtree rooted at t in sorted order.
14  */
15 void printTree( BinaryNode *t, ostream & out ) const
16 {
17     if( t != nullptr )
18     {
19         printTree( t->left, out );
20         out << t->element << endl;
21         printTree( t->right, out );
22     }
23 }

```

18

## Tree Traversals

- preorder traversal
  - visit node first, then left subtree, then right subtree
- postorder traversal
  - visit left subtree, right subtree, then node
- graphic technique for traversals
- level-order traversal
  - all nodes at depth  $d$  are processed before any node at depth  $d + 1$ 
    - not implemented with recursion
    - queue

19

## Tree Traversals

- height method using postorder traversal

```

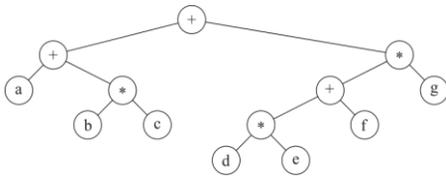
1  /**
2   * Internal method to compute the height of a subtree rooted at t.
3   */
4  int height( BinaryNode *t )
5  {
6      if( t == nullptr )
7          return -1;
8      else
9          return 1 + max( height( t->left ), height( t->right ) );
10 }

```

20

## Binary Tree Example: Expression Trees

- expression tree
  - leaves represent operands (constants or variable names)
  - interior nodes represent operators
  - binary tree since most operators are binary, but not required
  - some operations are unary

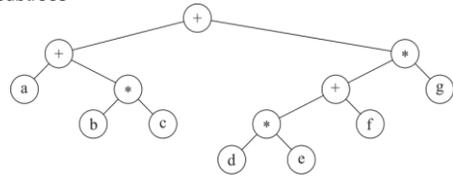


21

21

## Binary Tree Example: Expression Trees

- evaluate expression tree,  $T$ , by applying operator at root to values obtained by recursively evaluating left and right subtrees



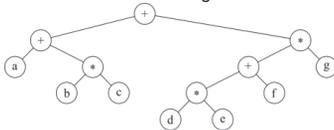
- left subtree:  $a + (b * c)$
- right subtree:  $((d * e) + f) * g$
- complete tree:  $(a + (b * c)) + (((d * e) + f) * g)$

22

22

## Binary Tree Example: Expression Trees

- inorder traversal
  - recursively produce left expression
  - print operator at root
  - recursively produce right expression
- postorder traversal
  - result:  $a b c * + d e * f + g * +$
- preorder traversal
  - result:  $++ a * b c * + * d e f g$



23

23

## Binary Tree Example: Expression Trees

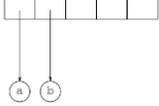
- goal: convert a postorder expression into an expression tree
  - read expression one symbol at a time
  - if operand, create node and push a pointer to it on the stack
  - if operator, pop pointers to two trees  $T_1$  and  $T_2$  from the stack
    - form new tree with operator as root
    - pointer to this tree is then pushed on the stack

24

24

## Binary Tree Example: Expression Trees

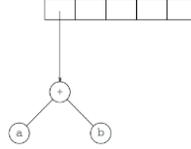
- example:  $a b + c d e + **$
- first two symbols are operands and are pushed on the stack



25

## Binary Tree Example: Expression Trees

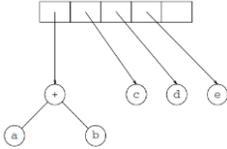
- example:  $a b + c d e + **$
- after + is read, two pointers are popped and new tree formed with a pointer pushed on the stack



26

## Binary Tree Example: Expression Trees

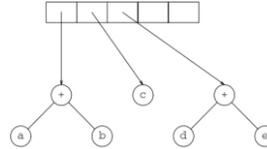
- example:  $a b + c d e + **$
- next, c, d, and e are read, with one-node tree created for each and pushed on the stack



27

## Binary Tree Example: Expression Trees

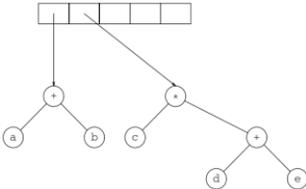
- example:  $a b + c d e + **$
- after + is read, two trees are merged



28

## Binary Tree Example: Expression Trees

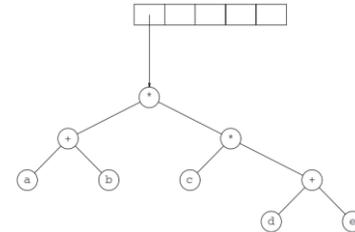
- example:  $a b + c d e + **$
- after \* is read, two trees are popped to form a new tree with a \* as root



29

## Binary Tree Example: Expression Trees

- example:  $a b + c d e + **$
- finally + is read, two trees are popped to form a final tree, which is left on the stack



30

25

26

27

28

29

30

## Binary Search Tree ADT

- binary trees often used for searching
- assume each node in the tree stores one element (integer)
- binary search tree
  - for every node  $X$  in the tree
    - all items in left subtree are smaller than  $X$
    - all items in right subtree are greater than  $X$
  - items in tree must be order-able



31

31

## Binary Search Tree ADT

- common operations on binary search trees
  - often written recursively
  - since average depth is  $O(\log N)$ , no worry about stack space
- binary search tree interface
  - searching depends on  $<$  operator, which must be defined for Comparable type
  - only data member is root pointer

32

32

## Binary Search Tree ADT

```

1  template <typename Comparable>
2  class BinarySearchTree
3  {
4  public:
5      BinarySearchTree();
6      BinarySearchTree(const BinarySearchTree & rhs);
7      BinarySearchTree(BinarySearchTree && rhs);
8      ~BinarySearchTree();
9
10     const Comparable & findMin() const;
11     const Comparable & findMax() const;
12     bool contains(const Comparable & x) const;
13     bool isEmpty() const;
14     void printTree(ostream & out = cout) const;
15
16     void makeEmpty();
17     void insert(const Comparable & x);
18     void insert(Comparable && x);
19     void remove(const Comparable & x);
20
21     BinarySearchTree & operator=(const BinarySearchTree & rhs);
22     BinarySearchTree & operator=(BinarySearchTree && rhs);

```

33

33

## Binary Search Tree ADT

```

24 private:
25     struct BinaryNode
26     {
27         Comparable element;
28         BinaryNode *left;
29         BinaryNode *right;
30
31         BinaryNode(const Comparable & theElement, BinaryNode *lt, BinaryNode *rt)
32             : element(theElement), left(lt), right(rt) {}
33
34         BinaryNode(Comparable && theElement, BinaryNode *lt, BinaryNode *rt)
35             : element(std::move(theElement)), left(lt), right(rt) {}
36     };
37
38     BinaryNode *root;
39
40     void insert(const Comparable & x, BinaryNode * & t);
41     void insert(Comparable && x, BinaryNode * & t);
42     void remove(const Comparable & x, BinaryNode * & t);
43     BinaryNode * findMin(BinaryNode * t) const;
44     BinaryNode * findMax(BinaryNode * t) const;
45     bool contains(const Comparable & x, BinaryNode * t) const;
46     void makeEmpty(BinaryNode * & t);
47     void printTree(BinaryNode * t, ostream & out) const;
48     BinaryNode * clone(BinaryNode * t) const;
49 };

```

34

34

## Binary Search Tree ADT

- test for item in subtree

```

1  /**
2   * Internal method to test if an item is in a subtree.
3   * x is item to search for.
4   * t is the node that roots the subtree.
5   */
6  bool contains(const Comparable & x, BinaryNode * t) const
7  {
8      if (t == nullptr)
9          return false;
10     else if (x == t->element)
11         return contains(x, t->left);
12     else if (t->element < x)
13         return contains(x, t->right);
14     else
15         return true; // Match
16 }

```

35

35

## Binary Search Tree ADT

- findMin** and **findMax**

- private methods return pointer to smallest/largest elements in the tree
- to find the minimum, start at the root and go left as long as possible
- similar for finding the maximum

36

36

## Binary Search Tree ADT

-recursive version of `findMin`

```

1  /**
2  * Internal method to find the smallest item in a subtree t.
3  * Return node containing the smallest item.
4  */
5  BinaryNode * findMin( BinaryNode *t ) const
6  {
7      if( t == nullptr )
8          return nullptr;
9      if( t->left == nullptr )
10         return t;
11     return findMin( t->left );
12 }

```

37

## Binary Search Tree ADT

-nonrecursive version of `findMax`

```

1  /**
2  * Internal method to find the largest item in a subtree t.
3  * Return node containing the largest item.
4  */
5  BinaryNode * findMax( BinaryNode *t ) const
6  {
7      if( t != nullptr )
8          while( t->right != nullptr )
9              t = t->right;
10     return t;
11 }

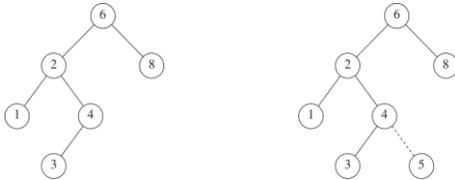
```

38

## Binary Search Tree ADT

## -insertion for binary search trees

- to insert  $X$  into tree  $T$ , proceed down the tree, as in the `contains` function
- if  $X$  is found, do nothing
- otherwise, insert  $X$  at the last spot on the path traversed
- example: insert 5 into binary search tree



39

39

## Binary Search Tree ADT

- duplicates can be handled by adding a count to the node record
- better than inserting duplicates in tree
- may not work well if key is only small part of larger structure

40

40

## Binary Search Tree ADT

## -deletion in binary search tree may be difficult

- multiple cases
- if node is leaf, it can be deleted immediately
- if node has only one child, node can be deleted after its parent adjusts a link to bypass the node



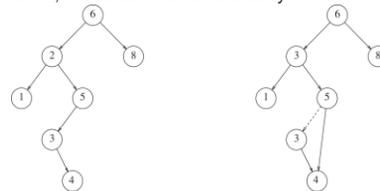
41

41

## Binary Search Tree ADT

## -multiple cases (cont.)

- complicated case: node with two children
- replace data of this node with smallest data of right subtree and recursively delete the node
- since smallest node in right subtree cannot have a left child, the second remove is easy



42

42

### Binary Search Tree ADT

- if number of deletions small, lazy deletion may be used
  - node is marked deleted rather than actually being deleted
- especially popular if duplicates allowed
  - count of duplicates can be decremented
- incurs only small penalty on tree since height not affected greatly
- if deleted node reinstated, some benefits

43

43

### Binary Search Tree Average-Case Analysis

- we expect most operations on binary search trees will have  $O(\log N)$  time
  - average depth over all nodes can be shown to be  $O(\log N)$
- all insertions and deletions must be equally likely
- sum of the depths of all nodes in a tree is known as internal path length

44

44

### Binary Search Tree Average-Case Analysis

- the run time of binary search trees depends on the depth of the tree, which in turn depends on the order that the keys are inserted
- let  $D(N)$  be the internal path length for a tree of  $N$  nodes
- we know that  $D(1) = 0$
- a tree of an  $i$ -node left subtree and an  $(N - i - 1)$ -node right subtree, plus a root at depth zero for  $0 \leq i \leq N$
- total number of nodes in tree = left subtree + right subtree + 1
- all nodes except the root are one level deeper,

$$D(N) = D(i) + D(N - i - 1) + N - 1$$

45

45

### Binary Search Tree Average-Case Analysis

- if all subtree sizes are equally likely, then the average for each subtree is

$$\left(\frac{1}{N}\right) \sum_{j=1}^{N-1} D(j)$$

therefore, for the total number of nodes

$$D(N) = \left(\frac{2}{N}\right) \left[ \sum_{j=1}^{N-1} D(j) \right] + N - 1$$

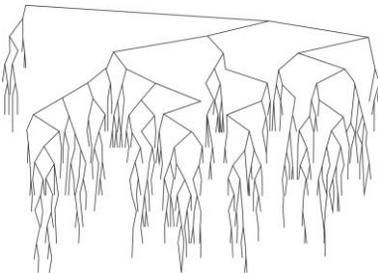
- once this recurrence relation is evaluated, the result is  $D(N) = O(N \log N)$  and the average number of nodes is  $O(\log N)$

46

46

### Binary Search Tree Average-Case Analysis

- example: randomly generated 500-node tree has expected depth of 9.98



47

47

### Binary Search Tree Average-Case Analysis

- deletions, however, bias the left subtrees to be longer because we always replace a deleted node with a node from the right subtree
- exact effect of deletions still unknown
- if insertions and deletions are alternated  $\theta(N^2)$  times, then expected depth is  $\theta(\sqrt{N})$

48

48



### AVL Trees

- minimum number of nodes,  $S(h)$ , in an AVL tree of height  $h$

$$S(h) = S(h - 1) + S(h - 2) + 1 \quad S(0) = 1, \quad S(1) = 2$$

- closely related to Fibonacci numbers
- all operations can be performed in  $O(\log N)$  time, except insertion and deletion

55

### AVL Trees

- insertion
  - update all balance information in the nodes on the path back to the root
  - could violate the balance condition
  - rotations used to restore the balance property
- deletion
  - perform same promotion as in a binary search tree, updating the balance information as necessary
  - same balancing operations for insertion can then be used

56

### AVL Trees

- if  $\alpha$  is the node requiring rebalancing (the heights of its left and right subtrees differ by 2), the violation occurred in one of four cases
  - an insertion into the left subtree of the left child of  $\alpha$
  - an insertion into the right subtree of the left child of  $\alpha$
  - an insertion into the left subtree of the right child of  $\alpha$
  - an insertion into the right subtree of the right child of  $\alpha$
- cases 1 and 4 are mirror image symmetries with respect to  $\alpha$  and can be resolved with a single rotation
- cases 2 and 3 are mirror image symmetries with respect to  $\alpha$  and can be resolved with a double rotation

57

### AVL Trees

- single rotation
  - only possible case 1 scenario
  - to balance, imagine "picking up" tree by  $k_1$

- new tree has same height as original tree

58

### AVL Trees

- example tree
  - when adding 6, node 8 becomes unbalanced
  - to balance, perform single rotation between 7 and 8

59

### AVL Trees

- example tree
  - symmetric case for case 4

60

**AVL Trees**

-example  
 -insert 3, 2, and 1 into an empty tree

61

**AVL Trees**

-example  
 -insert 4 and 5

62

**AVL Trees**

-example  
 -insert 6

63

**AVL Trees**

-example  
 -insert 7

64

**AVL Trees**

-double rotation  
 -for cases 2 and 3, a single rotation will not work

-tree Y can be expanded to a node with two subtrees

65

**AVL Trees**

-double rotation  
 -left-right double rotation for case 2

-right-left double rotation for case 3

66

### AVL Trees

-example  
-insert 16 and 15

67

### AVL Trees

-example  
-insert 14

68

### AVL Trees

-example  
-insert 13

69

### AVL Trees

-example  
-insert 12

70

### AVL Trees

-example  
-insert 11, 10, and 8

71

### AVL Trees

-example  
-insert 9

72

## AVL Trees

- implementation
- node definition

```

1 struct AvlNode
2 {
3     Comparable element;
4     AvlNode *left;
5     AvlNode *right;
6     int height;
7
8     AvlNode( const Comparable & ele, AvlNode *lt, AvlNode *rt, int h = 0 )
9         : element( ele ), left( lt ), right( rt ), height( h ) {}
10
11     AvlNode( Comparable && ele, AvlNode *lt, AvlNode *rt, int h = 0 )
12         : element( std::move( ele ) ), left( lt ), right( rt ), height( h ) {}
13 };

```

73

## AVL Trees

- implementation
- function to compute height of AVL node

```

1 /**
2  * Return the height of node t or -1 if nullptr.
3  */
4 int height( AvlNode *t ) const
5 {
6     return t == nullptr ? -1 : t->height;
7 }

```

74

## AVL Trees

- implementation
- insertion

```

1 /**
2  * Internal method to insert into a subtree.
3  * x is the item to insert.
4  * t is the node that roots the subtree.
5  * Set the new root of the subtree.
6  */
7 void insert( const Comparable &x, AvlNode * &t )
8 {
9     if( t == nullptr )
10         t = new AvlNode( x, nullptr, nullptr );
11     else if( x < t->element )
12         insert( x, t->left );
13     else if( t->element < x )
14         insert( x, t->right );
15     balance( t );
16 }

```

75

75

## AVL Trees

- implementation

```

19 static const int ALLOWED_IMBALANCE = 1;
20
21 // Assume t is balanced or within one of being balanced
22 void balance( AvlNode * &t )
23 {
24     if( t == nullptr )
25         return;
26
27     if( height( t->left ) - height( t->right ) > ALLOWED_IMBALANCE )
28         if( height( t->left->left ) >= height( t->left->right ) )
29             rotateWithLeftChild( t );
30         else
31             doubleWithLeftChild( t );
32     else
33         if( height( t->right ) - height( t->left ) > ALLOWED_IMBALANCE )
34             if( height( t->right->right ) >= height( t->right->left ) )
35                 rotateWithRightChild( t );
36             else
37                 doubleWithRightChild( t );
38     t->height = max( height( t->left ), height( t->right ) ) + 1;
39 }
40

```

76

76

## AVL Trees

- implementation
- single rotation

```

1 /**
2  * Rotate binary tree node with left child.
3  * For AVL trees, this is a single rotation for case 1.
4  * Update heights, then set new root.
5  */
6 void rotateWithLeftChild( AvlNode * &k2 )
7 {
8     AvlNode *k1 = k2->left;
9     k2->left = k1->right;
10    k1->right = k2;
11    k2->height = max( height( k2->left ), height( k2->right ) ) + 1;
12    k1->height = max( height( k1->left ), k2->height ) + 1;
13    k2 = k1;
14 }

```



77

77

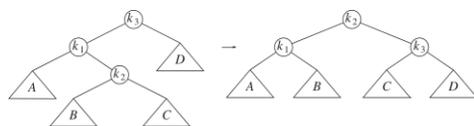
## AVL Trees

- implementation
- double rotation

```

1 /**
2  * Double rotate binary tree node: first left child
3  * with its right child; then node k3 with new left child.
4  * For AVL trees, this is a double rotation for case 2.
5  * Update heights, then set new root.
6  */
7 void doubleWithLeftChild( AvlNode * &k3 )
8 {
9     rotateWithRightChild( k3->left );
10    rotateWithLeftChild( k3 );
11 }

```



78

78

### AVL Trees

-implementation

-deletion

```

1 //**
2 * Internal method to remove from a subtree.
3 * x is the item to remove.
4 * t is the node that roots the subtree.
5 * Set the new root of the subtree.
6 */
7 void remove(const Comparable & x, AVLNode & t)
8 {
9     if (t == nullptr)
10        return; // item not found, do nothing
11
12     if (x < t->element)
13        remove(x, t->left);
14     else if (t->element < x)
15        remove(x, t->right);
16     else if (t->left != nullptr && t->right != nullptr) // Two children
17     {
18         t->element = frag(t->right, t->element);
19         remove(t->element, t->right);
20     }
21     else
22     {
23         AVLNode *oldNode = t;
24         t = (t->left != nullptr) ? t->left : t->right;
25         delete oldNode;
26     }
27
28     balance(t);
29 }
    
```

79

### Splay Trees

- different approach to ensuring  $O(\log N)$  behavior for tree operations (searches, insertions, and deletions).
- best case
  - splay tree operations may take  $N$  time
- however, splay trees make slow operations infrequently
- guarantee that  $M$  consecutive operations (insertions or deletions) requires at most  $O(M \log N)$ , so, on average, operations are  $O(\log N)$
- $O(\log N)$  is an amortized complexity
  - derivation is complex
- common for binary search trees to have a sustained sequence of bad accesses

80

### Splay Trees

- basic idea: when a node is accessed, it is moved to the top of the tree, with the thought that we might want to revisit recently accessed nodes more frequently
- use double rotations similar to AVL to move nodes to top of tree
- along the way, more branching is introduced in the tree, which reduces the height of the tree and thus the cost of tree operations

81

### Splay Trees

- single rotations don't work
- access  $k_i$

82

### Splay Trees

83

### Splay Trees

84

### Splay Trees

- double rotations consider parent and grandparent of accessed node
- zig: single branch (in one direction)
- zag: secondary branch (in opposite direction)
- when the parent node is the root, a single rotation for the zig is sufficient

85

### Splay Trees

- access X
- zig-zag

86

### Splay Trees

- access X
- zig-zig

87

### Splay Trees

- consider tree from previous example
- access  $k_1$
- zig-zag

88

### Splay Trees

- consider tree from previous example
- access  $k_1$
- zig-zig

89

### Splay Trees

- consider tree from previous example
- access  $k_1$

- $k_1$  is now at the root
- final tree has halved the distance of most nodes on the access path to the root

90

### Splay Trees

-example 2  
-access 1

-tree starts as worst case and results in much better structure for performance

91

### Splay Trees

-example 3: tree with only left children – access 1

92

### Splay Trees

-example 3: tree with only left children – access 2

93

### Splay Trees

-example 3: tree with only left children – access 3

94

### Splay Trees

-example 3: tree with only left children – access 4

95

### Splay Trees

-example 3: tree with only left children – access 5

96



### Top-Down Splay Trees

- instead, perform rotations on initial access path
  - result is faster
  - uses extra space  $O(1)$
  - retains amortized time bound of  $O(\log N)$

103

### Top-Down Splay Trees

- suppose we wish to access key  $i$
- during the access and concurrent splaying operation, the tree is broken into three parts
  - a left tree, which contains all the keys from the original tree known at the time to be less than  $i$
  - a right tree, which contains all the keys from the original tree known at the time to be greater than  $i$
  - a middle tree, which consists of the subtree of the original tree rooted at the current node on the access path
- initially, the left and right trees are empty and the middle tree is the entire tree
- at each step we tack bits of the middle tree onto the left and right trees

104

### Top-Down Splay Trees

- rotations for zig, zig-zig, and zig-zag cases

105

### Top-Down Splay Trees

- zig-zag case can be simplified to just a zig since no rotations are performed
- instead of making Z the root, we make Y the root

- simplifies coding, but only descends one level
- requires more iterations

106

### Top-Down Splay Trees

- after final splaying

107

### Top-Down Splay Trees

- example: access 19

108

### Top-Down Splay Trees

-example: access 19

zig-zig

zig

109

### Top-Down Splay Trees

-example: access 19

reassemble

110

### Top-Down Splay Trees

- use a header to hold the roots of the left and right subtrees
- left pointer will contain root of right subtree
- right pointer will contain root of left subtree
- easy to reconstruct at end of splaying

111

### Top-Down Splay Trees

-example 2: access 14

empty tree

empty tree

112

### Top-Down Splay Trees

-example 2: access 14

-start at root and look down two nodes along path to 14

113

### Top-Down Splay Trees

-example 2: access 14

empty tree

114

113

114

### Top-Down Splay Trees

-example 2: access 14  
 -continuing down the tree, this is a zig-zag condition

empty tree

115

### Top-Down Splay Trees

-example 2: access 14  
 -tree is reconfigured

empty tree

116

### Top-Down Splay Trees

-example 2: access 14  
 -simple zig

117

### Top-Down Splay Trees

-example 2: access 14  
 -move accessed node to root and reassemble tree

118

### B-Trees

-B-trees were developed in the late 1960s by Rudolf Bayer and Edward McCreight:

R. Bayer and E. McCreight, *Organization and maintenance of large ordered indexes*, Acta Informatica vol. 1, no. 3 (1972), pp. 173-189.

-originally motivated by applications in databases  
 -B-trees shown here really B+ tree

119

### B-Trees

-thus far, we have frequently treated the key as if it were the data being stored, but that is rarely the case

-example: student records in Banner

- most effective search key is W&M ID (e.g., 930...) since it is unique
- the record (value) associated with each key contains much more information
  - Student Information
  - Student Academic Transcript
  - Student Active Registrations
  - Student Schedule
  - Student E-mail Address
  - Student Address and Phones ...

120

### B-Trees

- B-trees are particularly useful when we cannot fit all of our data in memory, but have to perform reads and writes from secondary storage (e.g., disk drives)
  - disk accesses incredibly expensive, relatively speaking
  - consider a disk drive that rotates at 7200 rpm
    - the rotational speed plays a role in retrieval time; for a 7200 rpm disk, each revolution takes  $60/7200 = 1/120$  s, or about 8.3 ms
  - a typical seek time (the time for the disk head to move to the location where data will be read or written) for 7200 rpm disks is around 9 ms
  - this means we can perform 100-120 random disk accesses per second

121

### B-Trees

- meanwhile, our CPU can perform  $> 1,000,000,000$  operations per second
- suppose we have a database with  $N = 10,000,000$  entries that we organize in a tree
  - in an AVL tree, a worst-case search requires  $1.44 \lg N \approx 33$  disk accesses
  - at 9 ms per access, this requires about 300 ms, so on average we can perform less than 4 searches per second
  - we would expect 1000 worst-case searches to take  $300,000 \text{ ms} = 300 \text{ s}$ , or about 5 minutes
  - in this application, search trees with height  $\lg N$  are still too high!

122

### B-Trees

- height can be reduced if we allow more branching
- binary search trees only allow 2-way branching
- example: 5-ary 31-node tree with height 3

123

### B-Trees

- B-tree of order  $M$  is an  $M$ -ary tree with the following properties
  1. data items are stored at leaves
  2. nonleaf nodes (internal nodes) store up to  $M - 1$  keys to guide the searching: key  $i$  represents the smallest key in subtree  $i + 1$
  3. root is either a leaf or has between two and  $M$  children
  4. all nonleaf nodes (except the root) have between  $\lceil M/2 \rceil$  and  $M$  children
  5. all leaves are at the same depth and have between  $\lceil L/2 \rceil$  and  $L$  data items, for some  $L$

124

### B-Trees

- examples
  - for  $M = 2$ , there are between  $\lceil 2/2 \rceil = 1$  to 2 children
  - for  $M = 3$ , there are between  $\lceil 3/2 \rceil = 2$  to 3 children
  - for  $M = 4$ , there are between  $\lceil 4/2 \rceil = 2$  to 4 children
  - for  $M = 5$ , there are between  $\lceil 5/2 \rceil = 3$  to 5 children
  - for  $M = 42$ , there are between  $\lceil 42/2 \rceil = 21$  to 42 children
- requiring nodes to be half full guarantees that the tree will not degenerate into a simple binary search tree

125

### B-Trees

- examples:  $M = 5$

- all nonleaf nodes have between 3 and 5 children (and thus between 2 and 4 keys)
- root could have just 2 children
- here  $L$  is also 5: each leaf has between 3 and 5 data items

126

### B-Trees

- choosing  $M$  and  $L$ 
  - each node will occupy a disk block, say 8192 bytes, so we choose  $M$  and  $L$  based on the size of the items being stored
  - suppose each key uses 32 bytes and a link to another node uses 8 bytes
  - a node in a B-tree of order  $M$  has  $M - 1$  keys and  $M$  links, so a node requires
 
$$32(M - 1) + 8M = 40M - 32 \text{ bytes}$$
  - we choose the largest  $M$  that will allow a node to fit in a block
 
$$M = \left\lfloor \frac{8192 + 32}{40} \right\rfloor = 205$$

127

### B-Trees

- choosing  $M$  and  $L$  (cont.)
  - if the values are each 256 bytes, then we can fit
 
$$L = \left\lfloor \frac{8192}{256} \right\rfloor = 32$$
 in a single block
  - each leaf has between 16 and 32 values, and each internal node branches in at least 103 ways
  - if there are 1,000,000,000 values to store, there are at most 62,500,000 leaves
  - the leaves would be, in the worst case, on level
 
$$1 + \log_{103} 62,500,000 = 5$$
 so we can find data in at most 5 disk access
  - a BST would have at least  $1 + \log_2 62,500,000 = 27$  levels!

128

### B-Trees

- insertion: easy case – insert 57
  - first, follow the search tree to the correct leaf (external node)
  - if there are fewer than  $L$  items in the leaf, insert in the correct location
    - cost: 1 disk access
- insert 55?

129

### B-Trees

- insertion: splitting a leaf – insert 55
  - if there are already  $L$  items in the leaf
    - add the new item, split the node in two, and update the links in the parent node
    - cost: 3 disk accesses (one for each new node and one for the update of the parent node)

130

### B-Trees

- insertion: splitting a leaf – insert 55 (cont.)
  - the splitting rule ensures we still have a B-tree: each new node has at least  $\lfloor L/2 \rfloor$  values (e.g., if  $L = 3$ , there are 2 values in one node and 1 in the other, and if  $L = 4$ , each new node has 2 keys)

131

### B-Trees

- insertion: splitting a parent – insert 40
  - what if the parent node already has all the child nodes it can possibly have?
    - split the parent node, and update its parent
    - repeat until we arrive at the root
  - if necessary, split the root into two nodes and create a new root with the two nodes as children
    - this is why the root is allowed as few as 2 children
    - thus, a B-tree grows at the root

132

### B-Trees

-insertion: splitting a parent – insert 40 (cont.)

The first diagram shows a B-tree with root node [41, 66, 87]. It has three children: [8, 18], [48, 51, 54, 57], and [92, 97]. Each child has three leaf nodes. The second diagram shows the root node [26, 41, 66, 87] after a merge of the first two children and insertion of 40 into the second leaf of the second child.

133

### B-Trees

-insertion: other techniques – insert 29

- put a child up for adoption if a neighbor has room
- here, move 32 to the next leaf
- modifies parent, but keeps nodes fuller and saves space in the long run

The first diagram shows a B-tree with root node [26, 41, 66, 87]. It has three children: [8, 18], [35, 38], and [48, 51, 54, 57]. Each child has three leaf nodes. The second diagram shows the root node [20, 41, 66, 87] after moving 32 to the second leaf of the second child and inserting 29 into the first leaf of the second child.

134

### B-Trees

-deletion: delete 99

- could bring leaf below minimum number of data items
- adopt neighboring item if neighbor not at minimum
- otherwise, combine with neighbor to form a full leaf
- process could make its way up to the root
- if root left with 1 child, remove root and make its child the new root of the tree

The first diagram shows a B-tree with root node [26, 41, 66, 87]. It has three children: [8, 18], [35, 38], and [48, 51, 54, 57]. Each child has three leaf nodes. The second diagram shows the root node [26, 41, 66, 83] after deleting 99 from the last leaf and merging the last two leaves of the third child.

135

### B-Trees

-deletion: delete 99 (cont.)

The first diagram shows a B-tree with root node [26, 41, 66, 87]. It has three children: [8, 18], [35, 38], and [48, 51, 54, 57]. Each child has three leaf nodes. The second diagram shows the root node [26, 41, 66, 83] after deleting 99 from the last leaf and merging the last two leaves of the third child.

136