

Search Trees: BSTs and B-Trees

David Kauchak
cs302
Spring 2012



Administrative

- HW grading



Number guessing game

- I'm thinking of a number between 1 and n
- You are trying to guess the answer
- For each guess, I'll tell you "correct", "higher" or "lower"
- Describe an algorithm that minimizes the number of guesses



Binary Search Trees

- BST – A binary tree where a parent's value is greater than all values in the left subtree and less than or equal to all the values in the right subtree

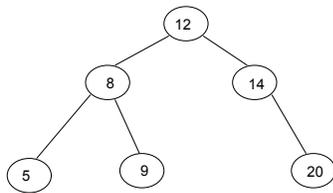
$$\text{leftTree}(i) < i \leq \text{rightTree}(i)$$

- the left and right children are also binary trees
- **Why not?**

$$\text{leftTree}(i) \leq i \leq \text{rightTree}(i)$$

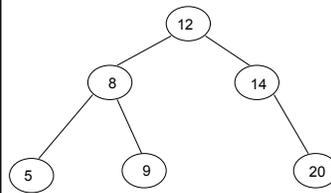
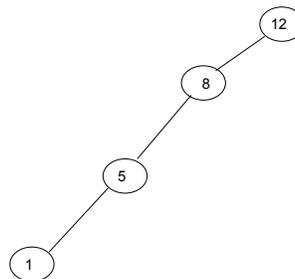
- Can be implemented with with pointers or an array



Example**What else can we say?**

$$\text{left}(i) < i \leq \text{right}(i)$$

- All elements to the left of a node are less than the node
- All elements to the right of a node are greater than or equal to the node
- The smallest element is the left-most element
- The largest element is the right-most element

**Another example: the loner****Another example: the twig**

Operations

- Search(T,k) – Does value k exist in tree T
- Insert(T,k) – Insert value k into tree T
- Delete(T,x) – Delete node x from tree T
- Minimum(T) – What is the smallest value in the tree?
- Maximum(T) – What is the largest value in the tree?
- Successor(T,x) – What is the next element in sorted order after x
- Predecessor(T,x) – What is the previous element in sorted order of x
- Median(T) – return the median of the values in tree T

Search

- How do we find an element?

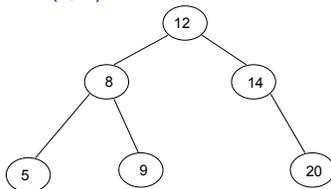
```

BSTSEARCH(x, k)
1  if x = null or k = x
2      return x
3  elseif k < x
4      return BSTSEARCH(LEFT(x), k)
5  else
6      return BSTSEARCH(RIGHT(x), k)
    
```

Finding an element

Search(T, 9)

$left(i) < i \leq right(i)$



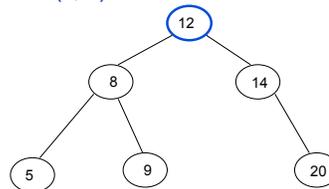
```

BSTSEARCH(x, k)
1  if x = null or k = x
2      return x
3  elseif k < x
4      return BSTSEARCH(LEFT(x), k)
5  else
6      return BSTSEARCH(RIGHT(x), k)
    
```

Finding an element

Search(T, 9)

$left(i) < i \leq right(i)$



```

BSTSEARCH(x, k)
1  if x = null or k = x
2      return x
3  elseif k < x
4      return BSTSEARCH(LEFT(x), k)
5  else
6      return BSTSEARCH(RIGHT(x), k)
    
```

Finding an element

$left(i) < i \leq right(i)$

Search(T, 9)

```

BSTSEARCH(x, k)
1 if x = null or k = x
2   return x
3 elseif k < x
4   return BSTSEARCH(LEFT(x), k)
5 else
6   return BSTSEARCH(RIGHT(x), k)
    
```

9 > 12?

Finding an element

$left(i) < i \leq right(i)$

Search(T, 9)

```

BSTSEARCH(x, k)
1 if x = null or k = x
2   return x
3 elseif k < x
4   return BSTSEARCH(LEFT(x), k)
5 else
6   return BSTSEARCH(RIGHT(x), k)
    
```

Finding an element

$left(i) < i \leq right(i)$

Search(T, 9)

```

BSTSEARCH(x, k)
1 if x = null or k = x
2   return x
3 elseif k < x
4   return BSTSEARCH(LEFT(x), k)
5 else
6   return BSTSEARCH(RIGHT(x), k)
    
```

Finding an element

$left(i) < i \leq right(i)$

Search(T, 13)

```

BSTSEARCH(x, k)
1 if x = null or k = x
2   return x
3 elseif k < x
4   return BSTSEARCH(LEFT(x), k)
5 else
6   return BSTSEARCH(RIGHT(x), k)
    
```

Finding an element

Search(T, 13) $left(i) < i \leq right(i)$

```

BSTSEARCH(x, k)
1  if x = null or k = x
2    return x
3  elseif k < x
4    return BSTSEARCH(LEFT(x), k)
5  else
6    return BSTSEARCH(RIGHT(x), k)
    
```

Finding an element

Search(T, 13) $left(i) < i \leq right(i)$

```

BSTSEARCH(x, k)
1  if x = null or k = x
2    return x
3  elseif k < x
4    return BSTSEARCH(LEFT(x), k)
5  else
6    return BSTSEARCH(RIGHT(x), k)
    
```

Finding an element

Search(T, 13) $left(i) < i \leq right(i)$

```

BSTSEARCH(x, k)
1  if x = null or k = x
2    return x
3  elseif k < x
4    return BSTSEARCH(LEFT(x), k)
5  else
6    return BSTSEARCH(RIGHT(x), k)
    
```

Iterative search

```

ITERATIVEBSTSEARCH(x, k)
1  while x ≠ null and k ≠ x
2    if k < x
3      x ← LEFT(x)
4    else
5      x ← RIGHT(x)
6  return x

BSTSEARCH(x, k)
1  if x = null or k = x
2    return x
3  elseif k < x
4    return BSTSEARCH(LEFT(x), k)
5  else
6    return BSTSEARCH(RIGHT(x), k)
    
```

Is BSTSearch correct?

```

BSTSEARCH(x, k)
1  if x = null or k = x
2      return x
3  elseif k < x
4      return BSTSEARCH(LEFT(x), k)
5  else
6      return BSTSEARCH(RIGHT(x), k)

```

$$\text{left}(i) < i \leq \text{right}(i)$$

Running time of BST

- **Worst case?**
 - $O(\text{height of the tree})$
- **Average case?**
 - $O(\text{height of the tree})$
- **Best case?**
 - $O(1)$

Height of the tree

- **Worst case height?**
 - $n-1$
 - “the twig”
- **Best case height?**
 - $\text{floor}(\log_2 n)$
 - complete (or near complete) binary tree
- **Average case height?**
 - Depends on two things:
 - the data
 - how we build the tree!

Insertion

```

BSTINSERT(T, x)
1  if ROOT(T) = null
2      ROOT(T) ← x
3  else
4      y ← ROOT(T)
5      while y ≠ null
6          prev ← y
7          if x < y
8              y ← LEFT(y)
9          else
10             y ← RIGHT(y)
11     PARENT(x) ← prev
12     if x < prev
13         LEFT(prev) ← x
14     else
15         RIGHT(prev) ← x

```

Insertion

```

BSTINSERT(T, x)
1  if ROOT(T) = null
2    ROOT(T) ← x
3  else
4    y ← ROOT(T)
5    while y ≠ null
6      prev ← y
7      if x < y
8        y ← LEFT(y)
9      else
10     y ← RIGHT(y)
11   PARENT(x) ← prev
12   if x < prev
13     LEFT(prev) ← x
14   else
15     RIGHT(prev) ← x
    
```

Similar to search

```

ITERATIVEBSTSEARCH(x, k)
1  while x ≠ null and k ≠ x
2    if k < x
3      x ← LEFT(x)
4    else
5      x ← RIGHT(x)
6  return x
    
```



Insertion

```

BSTINSERT(T, x)
1  if ROOT(T) = null
2    ROOT(T) ← x
3  else
4    y ← ROOT(T)
5    while y ≠ null
6      prev ← y
7      if x < y
8        y ← LEFT(y)
9      else
10     y ← RIGHT(y)
11   PARENT(x) ← prev
12   if x < prev
13     LEFT(prev) ← x
14   else
15     RIGHT(prev) ← x
    
```

Similar to search

Find the correct location in the tree



Insertion

```

BSTINSERT(T, x)
1  if ROOT(T) = null
2    ROOT(T) ← x
3  else
4    y ← ROOT(T)
5    while y ≠ null
6      prev ← y
7      if x < y
8        y ← LEFT(y)
9      else
10     y ← RIGHT(y)
11   PARENT(x) ← prev
12   if x < prev
13     LEFT(prev) ← x
14   else
15     RIGHT(prev) ← x
    
```

keeps track of the previous node we visited so when we fall off the tree, we know



Insertion

```

BSTINSERT(T, x)
1  if ROOT(T) = null
2    ROOT(T) ← x
3  else
4    y ← ROOT(T)
5    while y ≠ null
6      prev ← y
7      if x < y
8        y ← LEFT(y)
9      else
10     y ← RIGHT(y)
11   PARENT(x) ← prev
12   if x < prev
13     LEFT(prev) ← x
14   else
15     RIGHT(prev) ← x
    
```

add node onto the bottom of the tree



Correctness?

```

BSTINSERT(T, x)
1  if ROOT(T) = null
2    ROOT(T) ← x
3  else
4    y ← ROOT(T)
5    while y ≠ null
6      prev ← y
7      if x < y
8        y ← LEFT(y)
9      else
10     y ← RIGHT(y)
11   PARENT(x) ← prev
12   if x < prev
13     LEFT(prev) ← x
14   else
15     RIGHT(prev) ← x
    
```

maintain BST property

Correctness

```

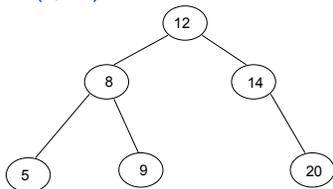
BSTINSERT(T, x)
1  if ROOT(T) = null
2    ROOT(T) ← x
3  else
4    y ← ROOT(T)
5    while y ≠ null
6      prev ← y
7      if x < y
8        y ← LEFT(y)
9      else
10     y ← RIGHT(y)
11   PARENT(x) ← prev
12   if x < prev
13     LEFT(prev) ← x
14   else
15     RIGHT(prev) ← x
    
```

What happens if it is a duplicate?

Inserting duplicate

Insert(T, 14)

$$left(i) < i \leq right(i)$$



Running time

```

BSTINSERT(T, x)
1  if ROOT(T) = null
2    ROOT(T) ← x
3  else
4    y ← ROOT(T)
5    while y ≠ null
6      prev ← y
7      if x < y
8        y ← LEFT(y)
9      else
10     y ← RIGHT(y)
11   PARENT(x) ← prev
12   if x < prev
13     LEFT(prev) ← x
14   else
15     RIGHT(prev) ← x
    
```

O(height of the tree)

Running time

```

BSTINSERT(T, x)
1  if ROOT(T) = null
2     ROOT(T) ← x
3  else
4     y ← ROOT(T)
5     while y ≠ null
6         prev ← y
7         if x < y
8             y ← LEFT(y)
9         else
10            y ← RIGHT(y)
11    PARENT(x) ← prev
12    if x < prev
13        LEFT(prev) ← x
14    else
15        RIGHT(prev) ← x

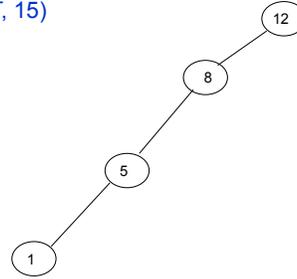
```

$O(\text{height of the tree})$

Why not
 $\Theta(\text{height of the tree})?$

Running time

Insert(T, 15)



Height of the tree

- Worst case: “the twig” – When will this happen?

```

BSTINSERT(T, x)
1  if ROOT(T) = null
2     ROOT(T) ← x
3  else
4     y ← ROOT(T)
5     while y ≠ null
6         prev ← y
7         if x < y
8             y ← LEFT(y)
9         else
10            y ← RIGHT(y)
11    PARENT(x) ← prev
12    if x < prev
13        LEFT(prev) ← x
14    else
15        RIGHT(prev) ← x

```

Height of the tree

- Best case: “complete” – When will this happen?

```

BSTINSERT(T, x)
1  if ROOT(T) = null
2     ROOT(T) ← x
3  else
4     y ← ROOT(T)
5     while y ≠ null
6         prev ← y
7         if x < y
8             y ← LEFT(y)
9         else
10            y ← RIGHT(y)
11    PARENT(x) ← prev
12    if x < prev
13        LEFT(prev) ← x
14    else
15        RIGHT(prev) ← x

```

Height of the tree

- Average case for random data?

```

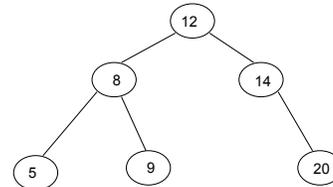
BSTINSERT(T, x)
1  if ROOT(T) = null
2    ROOT(T) ← x
3  else
4    y ← ROOT(T)
5    while y ≠ null
6      prev ← y
7      if x < y
8        y ← LEFT(y)
9      else
10     y ← RIGHT(y)
11  PARENT(x) ← prev
12  if x < prev
13    LEFT(prev) ← x
14  else
15    RIGHT(prev) ← x

```

Randomly inserted data into a BST generates a tree on average that is $O(\log n)$

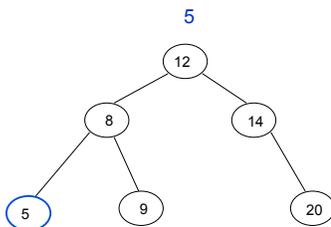
Visiting all nodes

- In sorted order



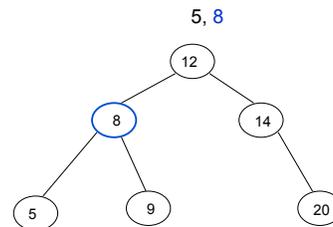
Visiting all nodes

- In sorted order



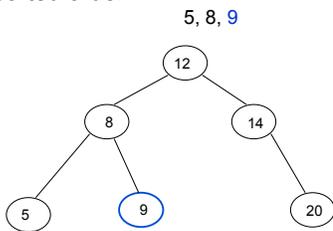
Visiting all nodes

- In sorted order



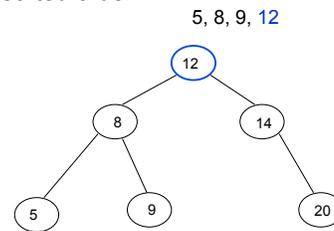
Visiting all nodes

- In sorted order



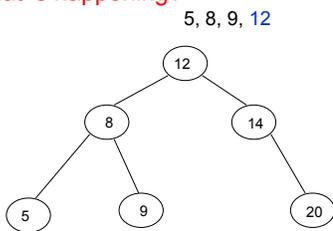
Visiting all nodes

- In sorted order



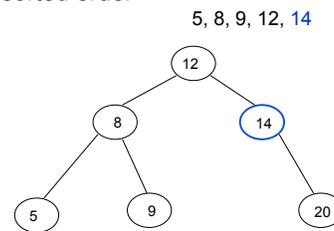
Visiting all nodes

- What's happening?



Visiting all nodes

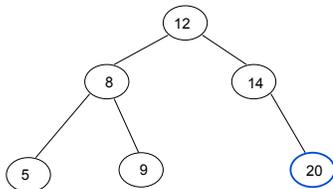
- In sorted order



Visiting all nodes

- In sorted order

5, 8, 9, 12, 14, 20



Visiting all nodes in order

```

INORDERTREEWALK(x)
1  if x ≠ null
2      INORDERTREEWALK(LEFT(x))
3      print x
4      INORDERTREEWALK(RIGHT(x))
  
```

Visiting all nodes in order

```

INORDERTREEWALK(x)
1  if x ≠ null
2      INORDERTREEWALK(LEFT(x))
3  print x
4      INORDERTREEWALK(RIGHT(x))
  
```

any operation

Is it correct?

```

INORDERTREEWALK(x)
1  if x ≠ null
2      INORDERTREEWALK(LEFT(x))
3      print x
4      INORDERTREEWALK(RIGHT(x))
  
```

- Does it print out all of the nodes in sorted order?

$$left(i) < i \leq right(i)$$

Running time?

```

INORDERTREEWALK(x)
1  if x ≠ null
2      INORDERTREEWALK(LEFT(x))
3      print x
4      INORDERTREEWALK(RIGHT(x))

```

- Recurrence relation:
 - j nodes in the left subtree
 - $n - j - 1$ in the right subtree

$$T(n) = T(j) + T(n - j - 1) + \Theta(1)$$

- Or
 - How much work is done for each call?
 - How many calls?
 - $\Theta(n)$

What about?

```

TREETWALK(x)
1  if x ≠ null
2      print x
3      TREETWALK(LEFT(x))
4      TREETWALK(RIGHT(x))

```

Preorder traversal

```

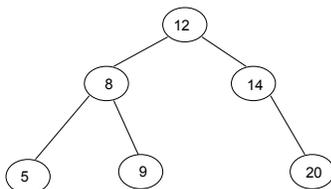
TREETWALK(x)
1  if x ≠ null
2      print x
3      TREETWALK(LEFT(x))
4      TREETWALK(RIGHT(x))

```

12, 8, 5, 9, 14, 20

How is this useful?

- Tree copying: insert in to new tree in preorder
- prefix notation: $(2+3)^*4 \rightarrow * + 2 3 4$



What about?

```

TREETWALK(x)
1  if x ≠ null
2      TREETWALK(LEFT(x))
3      TREETWALK(RIGHT(x))
4      print x

```

Postorder traversal

```

TREEWALK(x)
1 if x ≠ null
2   TREEWALK(LEFT(x))
3   TREEWALK(RIGHT(x))
4   print x
    
```

5, 9, 8, 20, 14, 12

How is this useful?

- postfix notation: $(2+3)*4 \rightarrow 4\ 3\ 2\ +\ *$
- ?

Min/Max

```

BSTMIN(x)
1 if LEFT(x) = null
2   return x
3 else
4   return BSTMIN(LEFT(x))
    
```

```

ITERATIVEBSTMIN(x)
1 while LEFT(x) ≠ null
2   x ← LEFT(x)
3 return x
    
```

Running time of min/max?

```

BSTMIN(x)
1 if LEFT(x) = null
2   return x
3 else
4   return BSTMIN(LEFT(x))
    
```

```

ITERATIVEBSTMIN(x)
1 while LEFT(x) ≠ null
2   x ← LEFT(x)
3 return x
    
```

$O(\text{height of the tree})$

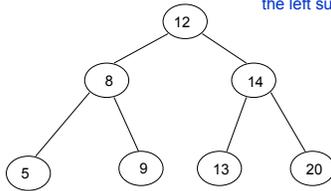
Successor and predecessor

Predecessor(12)? 9

Successor and predecessor

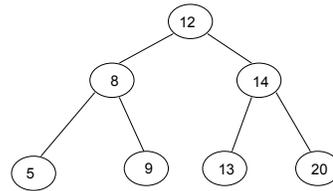
Predecessor in general? largest node of all those smaller than this node

rightmost element of the left subtree



Successor

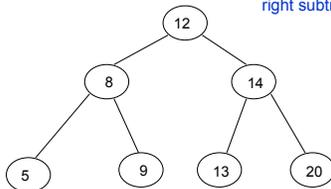
Successor(12)? 13



Successor

Successor in general? smallest node of all those larger than this node

leftmost element of the right subtree

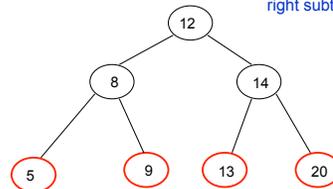


Successor

What if the node doesn't have a right subtree?

smallest node of all those larger than this node

leftmost element of the right subtree



Successor

What if the node doesn't have a right subtree?

- node is the largest
- the successor is the node that has x as a predecessor

Successor

successor is the node that has x as a predecessor

Successor

successor is the node that has x as a predecessor

Successor

• successor is the node that has x as a predecessor

Successor

keep going up until we're no longer a right child

- successor is the node that has x as a predecessor

Successor

```

SUCCESSOR(x)
1 if RIGHT(x) ≠ null
2   return BSTMIN(RIGHT(x))
3 else
4   y ← PARENT(x)
5   while y ≠ null and x = RIGHT(y)
6     x ← y
7     y ← PARENT(y)
8 return y
    
```

Successor

```

SUCCESSOR(x)
1 if RIGHT(x) ≠ null
2   return BSTMIN(RIGHT(x))
3 else
4   y ← PARENT(x)
5   while y ≠ null and x = RIGHT(y)
6     x ← y
7     y ← PARENT(y)
8 return y
    
```

if we have a right subtree, return the smallest of the right subtree

Successor

```

SUCCESSOR(x)
1 if RIGHT(x) ≠ null
2   return BSTMIN(RIGHT(x))
3 else
4   y ← PARENT(x)
5   while y ≠ null and x = RIGHT(y)
6     x ← y
7     y ← PARENT(y)
8 return y
    
```

find the node that x is the predecessor of

keep going up until we're no longer a right child

Successor running time

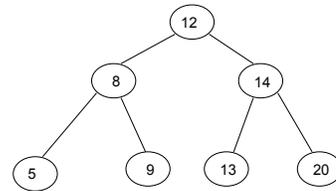
$O(\text{height of the tree})$

```

SUCCESSOR(x)
1  if RIGHT(x) ≠ null
2     return BSTMIN(RIGHT(x))
3  else
4     y ← PARENT(x)
5     while y ≠ null and x = RIGHT(y)
6         x ← y
7         y ← PARENT(y)
8  return y

```

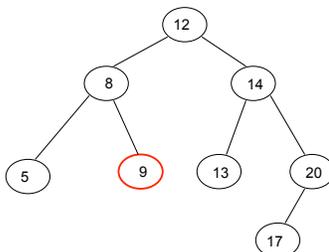
Deletion



Three cases!

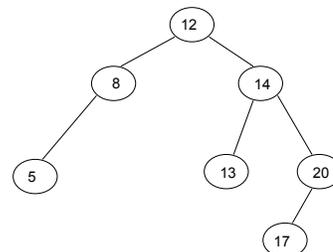
Deletion: case 1

- No children
- Just delete the node



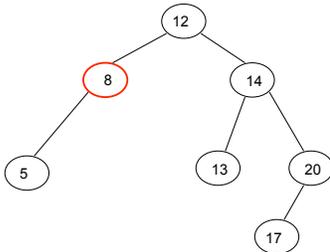
Deletion: case 1

- No children
- Just delete the node

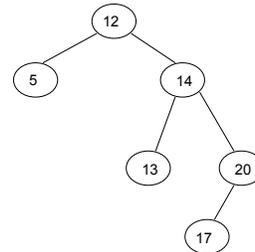


Deletion: case 2

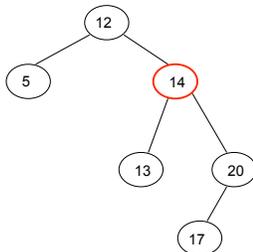
- One child
- Splice out the node

**Deletion: case 2**

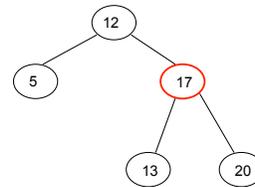
- One child
- Splice out the node

**Deletion: case 3**

- Two children
- Replace x with it's successor

**Deletion: case 3**

- Two children
- Replace x with it's successor



Deletion: case 3

- Two children
- Will we always have a successor?
- Why successor?
 - Case 1 or case 2 deletion
 - Larger than the left subtree
 - Less than or equal to right subtree

Height of the tree

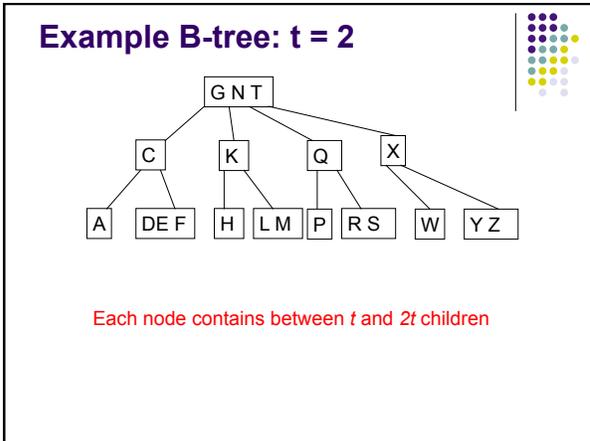
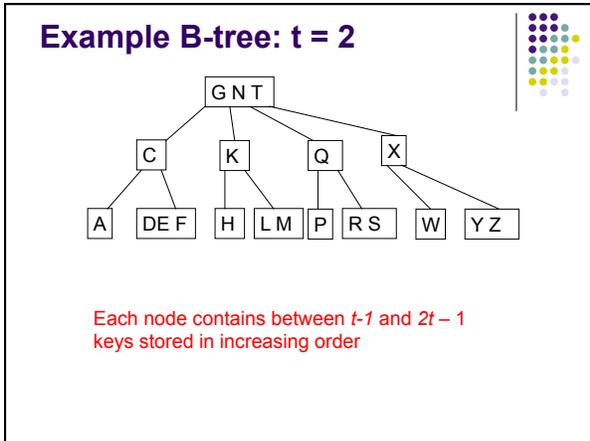
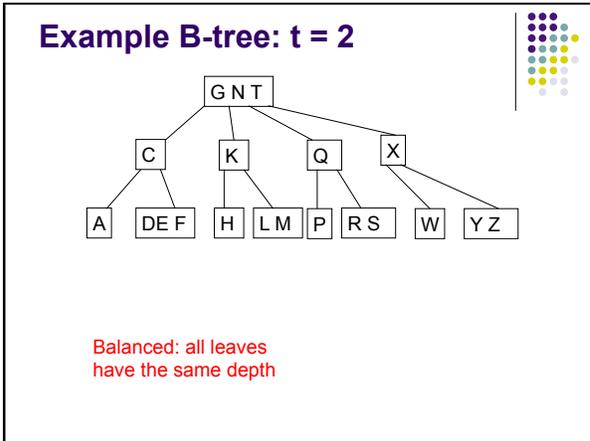
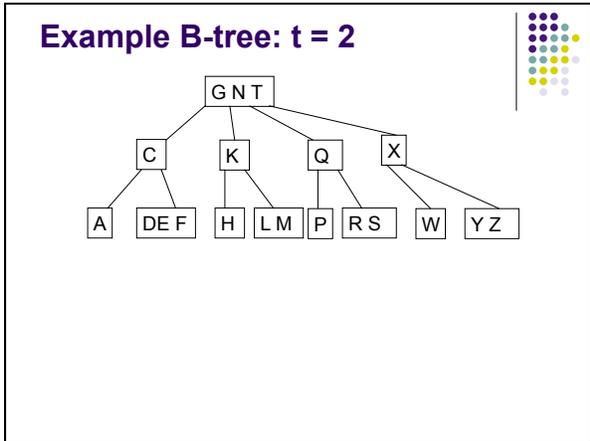
- Most of the operations take time $O(\text{height of the tree})$
- We said trees built from random data have height $O(\log n)$, which is asymptotically tight
- Two problems:
 - We can't always insure random data
 - What happens when we delete nodes and insert others after building a tree?

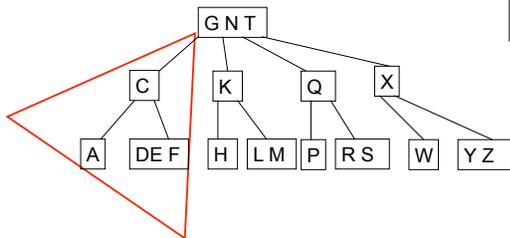
Balanced trees

- Make sure that the trees remain balanced!
 - Red-black trees
 - AVL trees
 - 2-3-4 trees
 - ...
- B-trees

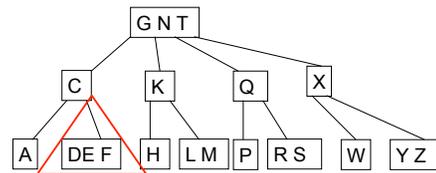
B-tree

- Defined by one parameter: t
- Balanced n -ary tree
- Each node contains between $t-1$ and $2t-1$ keys/data values (i.e. multiple data values per tree node)
 - keys/data are stored in **sorted order**
 - one exception: root can have $< t-1$ keys
- Each internal node contains between t and $2t$ children
 - the keys of a parent **delimit** the values of the children keys
 - For example, if $\text{key}_i = 15$ and $\text{key}_{i+1} = 25$ then child $i + 1$ must have keys between 15 and 25
- all leaves have the same depth

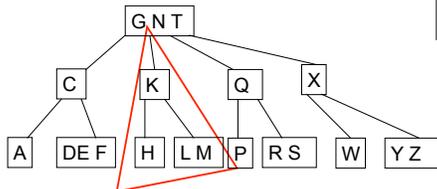


Example B-tree: $t = 2$ 

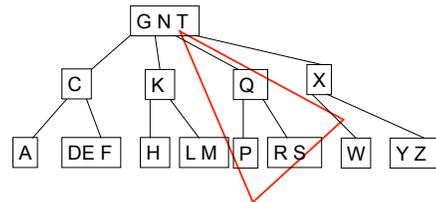
The keys of a parent delimit the values that a child's keys can take

Example B-tree: $t = 2$ 

The keys of a parent delimit the values that a child's keys can take

Example B-tree: $t = 2$ 

The keys of a parent delimit the values that a child's keys can take

Example B-tree: $t = 2$ 

The keys of a parent delimit the values that a child's keys can take

When do we use B-trees over other balanced trees?



- B-trees are generally an **on-disk** data structure
- Memory is limited or there is a large amount of data to be stored
- In the extreme, only one node is kept in memory and the rest on disk
- Size of the nodes is often determined by a page size on disk. *Why?*
- Databases frequently use B-trees

Notes about B-trees



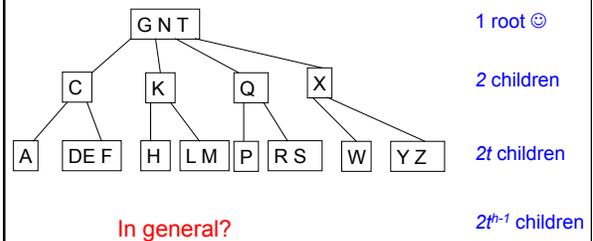
- Because t is generally large, the height of a B-tree is usually small
 - $t = 1001$ with height 2 can have over one billion values
- We will count both run-time as well as the number of disk accesses. *Why?*

Height of a B-tree



- B-trees have a similar feeling to BSTs
- We saw for BSTs that most of the operations depended on the height of the tree
- *How can we bound the height of the tree?*
- We know that nodes must have a minimum number of keys/data items
- *For a tree of height h , what is the smallest number of keys?*

Minimum number of nodes at each depth?



Minimum number of keys/values

$$n \geq 1 + (t-1) \sum_{i=1}^h 2t^{i-1}$$

Diagram labels: root (pointing to 1), min. keys per node (pointing to t-1), min. number of nodes (pointing to the summation term).

Minimum number of nodes

$$n \geq 1 + (t-1) \sum_{i=1}^h 2t^{i-1}$$

$$= 1 + 2(t-1) \left(\frac{t^h - 1}{t-1} \right)$$

$$= 2t^h - 1$$

so,

$$t^h \leq (n+1)/2$$

$$h \leq \log_t \frac{(n+1)}{2}$$

Searching B-Trees

Find value k into B-Tree node x

```

B-TREESEARCH(x, k)
1  i ← 1
2  while i ≤ n(x) and k > Kx[i]
3    i ← i + 1
4  if i ≤ n(x) and k = Kx[i]
5    return (x, i)
6  if LEAF(x)
7    return null
8  else
9    DISKREAD(Cx[i])
10   return B-TREESEARCH(Cx[i], k)
    
```

Diagram labels: number of keys (pointing to n(x)), key[i] (pointing to K_x[i]), child[i] (pointing to C_x[i]).

Searching B-Trees

```

B-TREESEARCH(x, k)
1  i ← 1
2  while i ≤ n(x) and k > Kx[i]
3    i ← i + 1
4  if i ≤ n(x) and k = Kx[i]
5    return (x, i)
6  if LEAF(x)
7    return null
8  else
9    DISKREAD(Cx[i])
10   return B-TREESEARCH(Cx[i], k)
    
```

Annotation: make disk reads explicit (pointing to line 9).

Searching B-Trees

B-TREESEARCH(x, k)

```

1   $i \leftarrow 1$ 
2  while  $i \leq n(x)$  and  $k > K_x[i]$ 
3       $i \leftarrow i + 1$ 
4  if  $i \leq n(x)$  and  $k = K_x[i]$ 
5      return  $(x, i)$ 
6  if LEAF( $x$ )
7      return null
8  else
9      DISKREAD( $C_x[i]$ )
10     return B-TREESEARCH( $C_x[i], k$ )

```

iterate through the sorted keys
and find the correct location

Searching B-Trees

B-TREESEARCH(x, k)

```

1   $i \leftarrow 1$ 
2  while  $i \leq n(x)$  and  $k > K_x[i]$ 
3       $i \leftarrow i + 1$ 
4  if  $i \leq n(x)$  and  $k = K_x[i]$ 
5      return  $(x, i)$ 
6  if LEAF( $x$ )
7      return null
8  else
9      DISKREAD( $C_x[i]$ )
10     return B-TREESEARCH( $C_x[i], k$ )

```

if we find the value
in this node, return it

Searching B-Trees

B-TREESEARCH(x, k)

```

1   $i \leftarrow 1$ 
2  while  $i \leq n(x)$  and  $k > K_x[i]$ 
3       $i \leftarrow i + 1$ 
4  if  $i \leq n(x)$  and  $k = K_x[i]$ 
5      return  $(x, i)$ 
6  if LEAF( $x$ )
7      return null
8  else
9      DISKREAD( $C_x[i]$ )
10     return B-TREESEARCH( $C_x[i], k$ )

```

if it's a leaf and we
didn't find it, it's not in
the tree

Searching B-Trees

B-TREESEARCH(x, k)

```

1   $i \leftarrow 1$ 
2  while  $i \leq n(x)$  and  $k > K_x[i]$ 
3       $i \leftarrow i + 1$ 
4  if  $i \leq n(x)$  and  $k = K_x[i]$ 
5      return  $(x, i)$ 
6  if LEAF( $x$ )
7      return null
8  else
9      DISKREAD( $C_x[i]$ )
10     return B-TREESEARCH( $C_x[i], k$ )

```

Recurse on the proper
child where the value is
between the keys

Search example: R

```

B-TREESEARCH(x, k)
1  i ← 1
2  while i ≤ n(x) and k > Kx[i]
3    i ← i + 1
4  if i ≤ n(x) and k = Kx[i]
5    return (x, i)
6  if LEAF(x)
7    return null
8  else
9    DISKREAD(Cx[i])
10   return B-TREESEARCH(Cx[i], k)
    
```

Search example: R

```

B-TREESEARCH(x, k)
1  i ← 1
2  while i ≤ n(x) and k > Kx[i]
3    i ← i + 1
4  if i ≤ n(x) and k = Kx[i]
5    return (x, i)
6  if LEAF(x)
7    return null
8  else
9    DISKREAD(Cx[i])
10   return B-TREESEARCH(Cx[i], k)
    
```

Search example: R

```

B-TREESEARCH(x, k)
1  i ← 1
2  while i ≤ n(x) and k > Kx[i]
3    i ← i + 1
4  if i ≤ n(x) and k = Kx[i]
5    return (x, i)
6  if LEAF(x)
7    return null
8  else
9    DISKREAD(Cx[i])
10   return B-TREESEARCH(Cx[i], k)
    
```

find the correct location

Search example: R

```

B-TREESEARCH(x, k)
1  i ← 1
2  while i ≤ n(x) and k > Kx[i]
3    i ← i + 1
4  if i ≤ n(x) and k = Kx[i]
5    return (x, i)
6  if LEAF(x)
7    return null
8  else
9    DISKREAD(Cx[i])
10   return B-TREESEARCH(Cx[i], k)
    
```

the value is not in this node

Search example: R

```

B-TREESearch(x, k)
1 i ← 1
2 while i ≤ n(x) and k > Kx[i]
3   i ← i + 1
4 if i ≤ n(x) and k = Kx[i]
5   return (x, i)
6 if LEAF(x)
7   return null
8 else
9   DISKREAD(Cx[i])
10  return B-TREESearch(Cx[i], k)
  
```

this is not a leaf node

Search example: R

```

B-TREESearch(x, k)
1 i ← 1
2 while i ≤ n(x) and k > Kx[i]
3   i ← i + 1
4 if i ≤ n(x) and k = Kx[i]
5   return (x, i)
6 if LEAF(x)
7   return null
8 else
9   DISKREAD(Cx[i])
10  return B-TREESearch(Cx[i], k)
  
```

Search example: R

```

B-TREESearch(x, k)
1 i ← 1
2 while i ≤ n(x) and k > Kx[i]
3   i ← i + 1
4 if i ≤ n(x) and k = Kx[i]
5   return (x, i)
6 if LEAF(x)
7   return null
8 else
9   DISKREAD(Cx[i])
10  return B-TREESearch(Cx[i], k)
  
```

find the correct location

Search example: R

```

B-TREESearch(x, k)
1 i ← 1
2 while i ≤ n(x) and k > Kx[i]
3   i ← i + 1
4 if i ≤ n(x) and k = Kx[i]
5   return (x, i)
6 if LEAF(x)
7   return null
8 else
9   DISKREAD(Cx[i])
10  return B-TREESearch(Cx[i], k)
  
```

not in this node and this is not a leaf

Search example: R

```

B-TREESEARCH(x, k)
1  i ← 1
2  while i ≤ n(x) and k > Kx[i]
3    i ← i + 1
4  if i ≤ n(x) and k = Kx[i]
5    return (x, i)
6  if LEAF(x)
7    return null
8  else
9    DISKREAD(Cx[i])
10   return B-TREESEARCH(Cx[i], k)
    
```

Search example: R

```

B-TREESEARCH(x, k)
1  i ← 1
2  while i ≤ n(x) and k > Kx[i]
3    i ← i + 1
4  if i ≤ n(x) and k = Kx[i]
5    return (x, i)
6  if LEAF(x)
7    return null
8  else
9    DISKREAD(Cx[i])
10   return B-TREESEARCH(Cx[i], k)
    
```

find the correct location

Search example: R

```

B-TREESEARCH(x, k)
1  i ← 1
2  while i ≤ n(x) and k > Kx[i]
3    i ← i + 1
4  if i ≤ n(x) and k = Kx[i]
5    return (x, i)
6  if LEAF(x)
7    return null
8  else
9    DISKREAD(Cx[i])
10   return B-TREESEARCH(Cx[i], k)
    
```

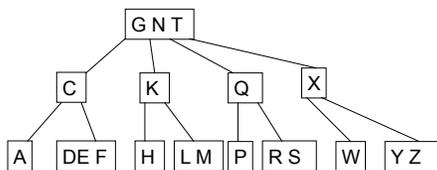
Search running time

- How many calls to BTreeSearch?
 - O(height of the tree)
 - O(log_n)
- Disk accesses?
 - One for each call – O(log_n)
- Computational time?
 - O(t) keys per node
 - linear search
 - O(t log_n)
- Why not binary search to find key in a node?

```

B-TREESEARCH(x, k)
1  i ← 1
2  while i ≤ n(x) and k > Kx[i]
3    i ← i + 1
4  if i ≤ n(x) and k = Kx[i]
5    return (x, i)
6  if LEAF(x)
7    return null
8  else
9    DISKREAD(Cx[i])
10   return B-TREESEARCH(Cx[i], k)
    
```

BST-Insert



B-Tree insert

- Starting at root, follow the *search* path down the tree
 - If the node is full (contains $2t - 1$ keys)
 - split the keys into two nodes around the median value
 - add the median value to the parent node
 - If the node is a leaf, insert it into the correct spot
- Observations
 - Insertions **always** happens in the leaves
 - When does the height of a B-tree grow?
 - Why do we know it's always ok when we're splitting a node to insert the median value into the parent?

Insertion: $t = 2$

GCNAHEKQMFWLTZDPRXYS

Insertion: $t = 2$

GCNAHEKQMFWLTZDPRXYS

G

Insertion: t = 2

GCNAHEKQMF^AWLTZDPRXYS

CG



Insertion: t = 2

GCNAHEKQMF^AWLTZDPRXYS

CGN



Insertion: t = 2

GCNAHEKQMF^AWLTZDPRXYS

CGN

Node is full, so split

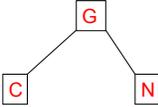


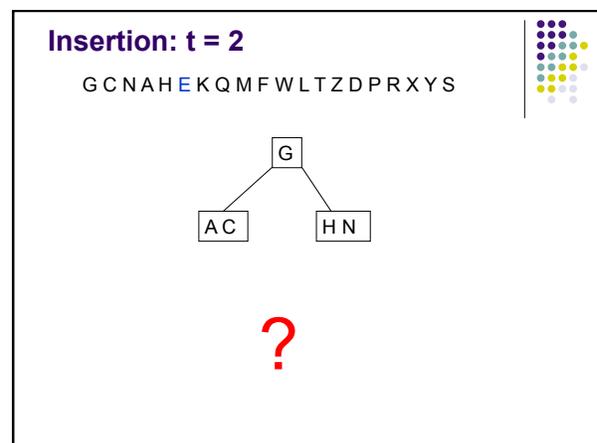
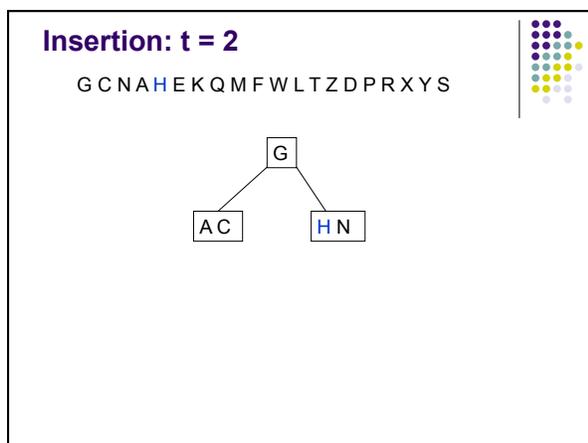
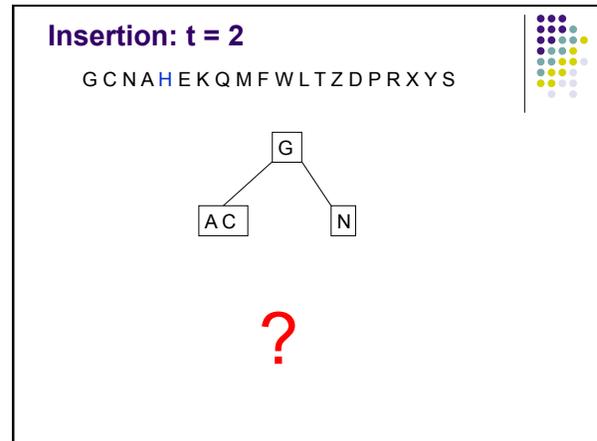
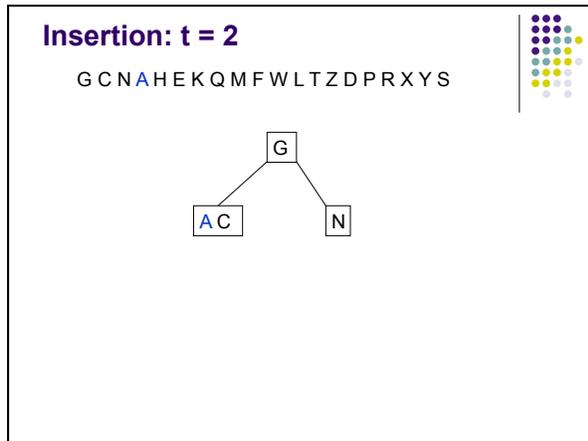
Insertion: t = 2

GCNAHEKQMF^AWLTZDPRXYS

CGN

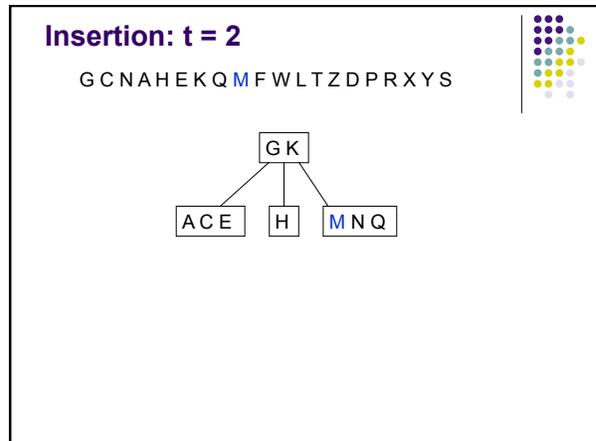
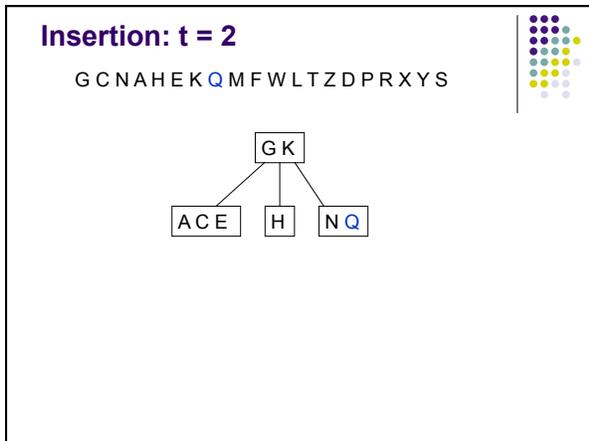
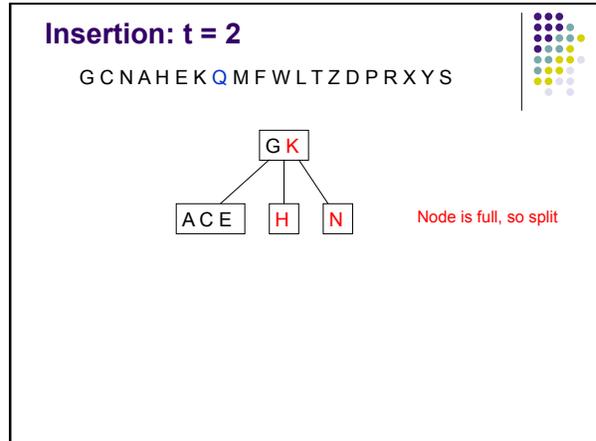
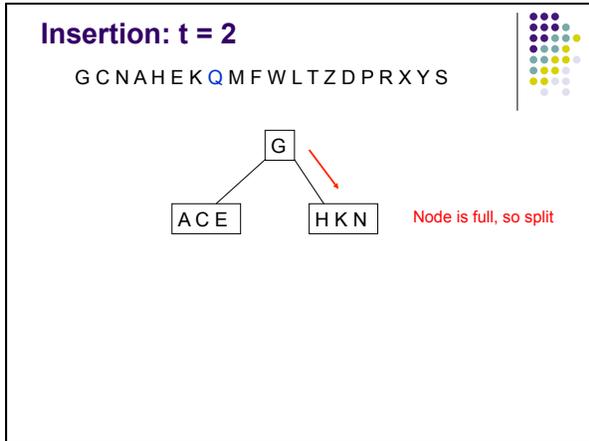
Node is full, so split





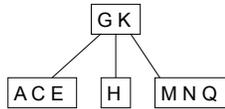
Insertion: t = 2

GCNAHEKQMFWLTZDPRXYS



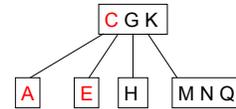
Insertion: $t = 2$

GCNAHEKQMFWLTZDPRXYS



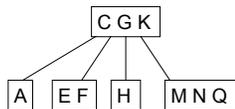
Insertion: $t = 2$

GCNAHEKQMFWLTZDPRXYS



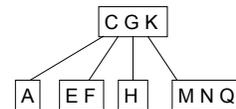
Insertion: $t = 2$

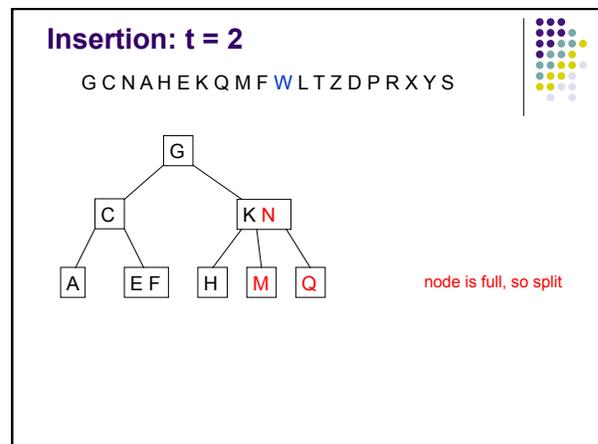
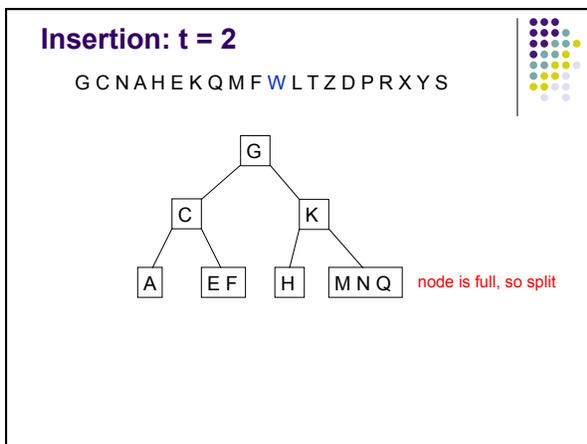
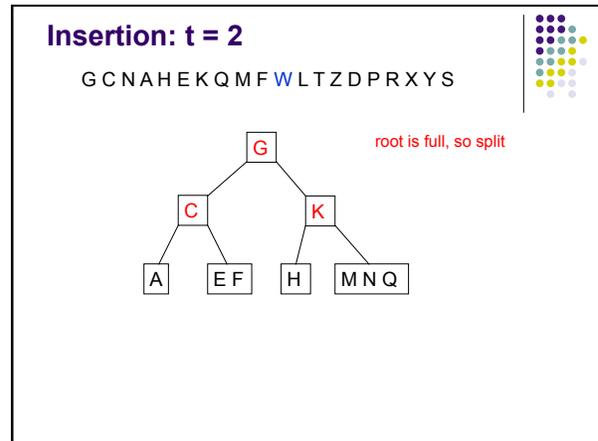
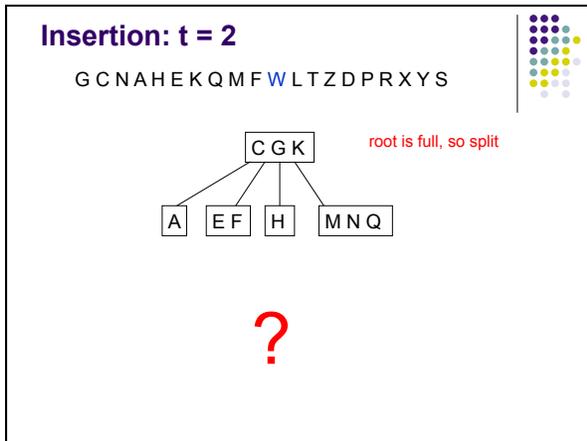
GCNAHEKQMFWLTZDPRXYS

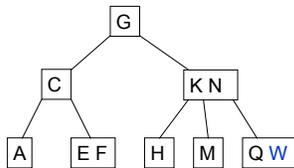
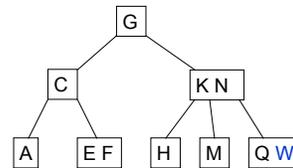


Insertion: $t = 2$

GCNAHEKQMFWLTZDPRXYS





Insertion: $t = 2$ GCNAHEKQMF **W**L T Z D P R X Y S**Insertion: $t = 2$** GCNAHEKQMF **W** ...**Correctness of insert**

- Starting at root, follow *search* path down the tree
 - If the node is full (contains $2t - 1$ keys), split the keys around the median value into two nodes and add the median value to the parent node
 - If the node is a leaf, insert it into the correct spot
- Does it add the value in the correct spot?
 - Follows the correct *search* path
 - Inserts in correct position

Correctness of insert

- Starting at root, follow *search* path down the tree
 - If the node is full (contains $2t - 1$ keys), split the keys around the median value into two nodes and add the median value to the parent node
 - If the node is a leaf, insert it into the correct spot
- Do we maintain a proper B-tree?
 - Maintain $t-1$ to $2t-1$ keys per node?
 - Always split full nodes when we see them
 - Only split full nodes
 - All leaves at the same level?
 - Only add nodes at leaves

Insert running time



- **Without any splitting?**
 - Similar to BTreeSearch, with one extra disk write at the leaf
 - $O(\log_t n)$ disk accesses
 - $O(t \log_t n)$ computation time

When a node is split



- **How many disk accesses?**
 - 3 disk write operations
 - 2 for the new nodes created by the split (one is reused, but must be updated)
 - 1 for the parent node to add median value
- **Runtime to split a node?**
 - $O(t)$ – iterating through the elements a few times since they're already in sorted order
- **Maximum number of nodes split for a call to insert?**
 - $O(\text{height of the tree})$

Running time of insert



- $O(\log_t n)$ disk accesses
- $O(t \log_t n)$ computational costs

Deleting a node from a B-tree



- Similar to insertion
 - must make sure we maintain B-tree properties (i.e. all leaves same depth and key/node restrictions)
 - Proactively move a key from a child to a parent if the parent has $t-1$ keys
- $O(\log_t n)$ disk accesses
- $O(t \log_t n)$ computational costs

Summary of operations



- Search, Insertion, Deletion
 - disk accesses: $O(\log n)$
 - computation: $O(t \log n)$
- Max, Min
 - disk accesses: $O(\log n)$
 - computation: $O(\log n)$
- Tree traversal
 - disk accesses: if $2t \sim$ page size: $O(\text{minimum \# pages to store data})$
 - Computation: $O(n)$