

Binary Search Trees

- ▶ binary search tree
- ▶ ordered operations
- ▶ deletion

References:
 Algorithms in Java, Chapter 12
<http://www.cs.princeton.edu/algs4/42bst>

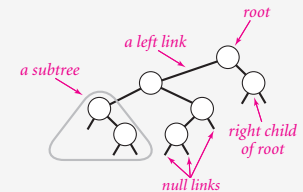
Algorithms in Java, 4th Edition · Robert Sedgewick and Kevin Wayne · Copyright © 2008 · October 17, 2008 8:03:06 AM

Binary search trees

Def. A BST is a binary tree in symmetric order.

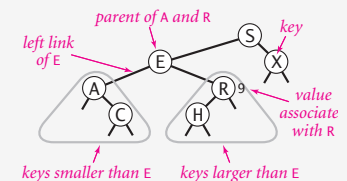
A binary tree is either:

- Empty.
- A key-value pair and two disjoint binary trees.



Symmetric order. Every node's key is:

- Larger than all keys in its left subtree.
- Smaller than all keys in its right subtree.



BST representation in Java

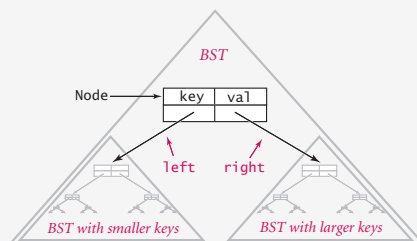
A **BST** is a reference to a root node.

A **Node** is comprised of four fields:

- A key and a value.
- A reference to the left and right subtree.

smaller keys larger keys

```
private class Node
{
    private Key key;
    private Value val;
    private Node left, right;
    public Node(Key key, Value val)
    {
        this.key = key;
        this.val = val;
    }
}
```



Key and Value are generic types; Key is Comparable

3

BST implementation (skeleton)

```
public class BST<Key extends Comparable<Key>, Value>
{
    private Node root;

    private class Node
    { /* see previous slide */ }

    public void put(Key key, Value val)
    { /* see next slides */ }

    public Value get(Key key)
    { /* see next slides */ }

    public void delete(Key key)
    { /* see next slides */ }

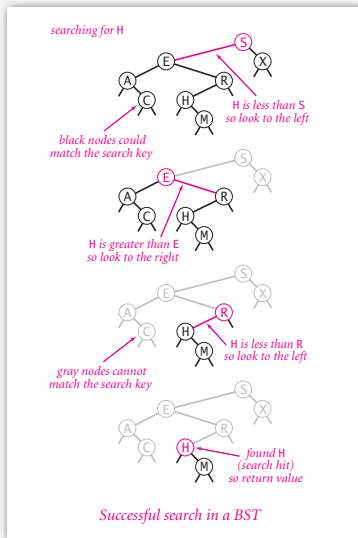
    public Iterable<Key> iterator()
    { /* see next slides */ }
}
```

root of BST

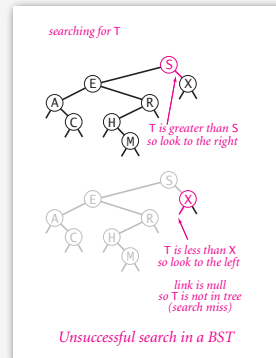
4

BST search

Get. Return value corresponding to given key, or `null` if no such key.



5



BST search: Java implementation

Get. Return value corresponding to given key, or `null` if no such key.

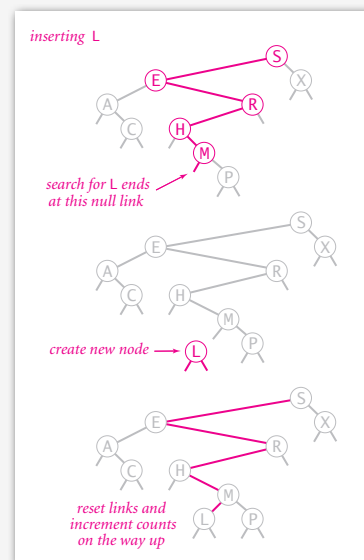
```
public Value get(Key key)
{
    Node x = root;
    while (x != null)
    {
        int cmp = key.compareTo(x.key);
        if (cmp < 0) x = x.left;
        else if (cmp > 0) x = x.right;
        else if (cmp == 0) return x.val;
    }
    return null;
}
```

Running time. Proportional to depth of node.

6

BST insert

Put. Associate value with key.



7

BST insert: Java implementation

Put. Associate value with key.

```
public void put(Key key, Value val)
{
    root = put(root, key, val);
}

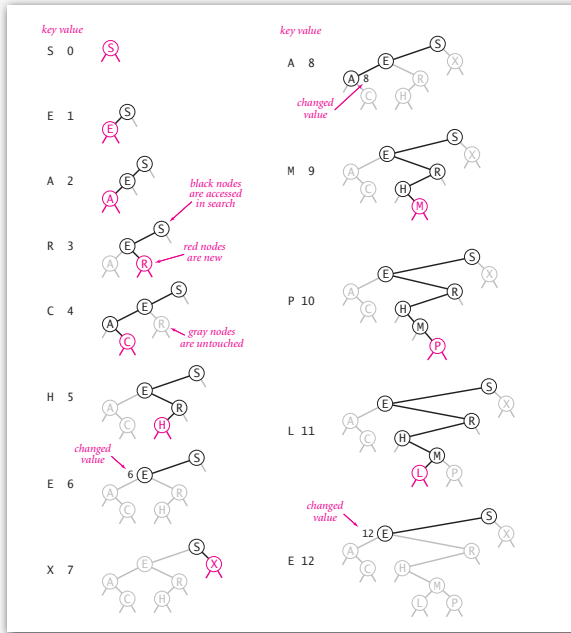
private Node put(Node x, Key key, Value val)
{
    if (x == null) return new Node(key, val);
    int cmp = key.compareTo(x.key);
    if (cmp < 0) x.left = put(x.left, key, val);
    else if (cmp > 0) x.right = put(x.right, key, val);
    else if (cmp == 0) x.val = val;
    return x;
}
```

← concise, but tricky, recursive code; read carefully!

Running time. Proportional to depth of node.

8

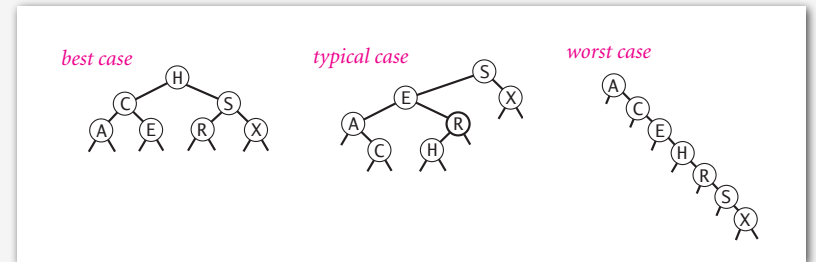
BST trace: standard indexing client



9

Tree shape

- Many BSTs correspond to same set of keys.
- Cost of search/insert is proportional to depth of node.

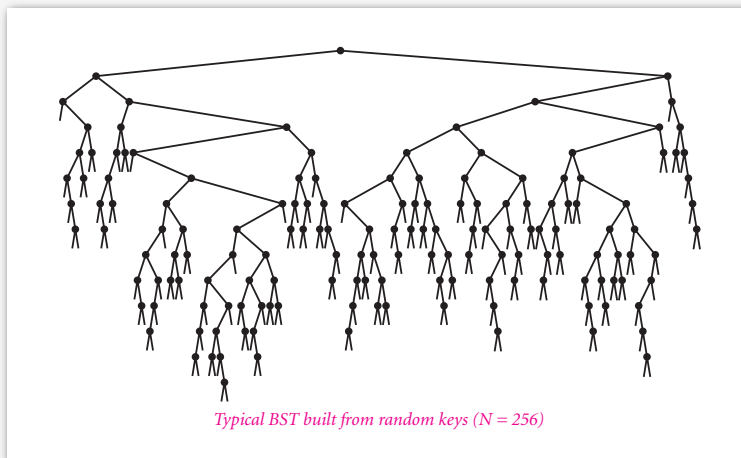


Remark. Tree shape depends on order of insertion.

10

BST insertion: random order

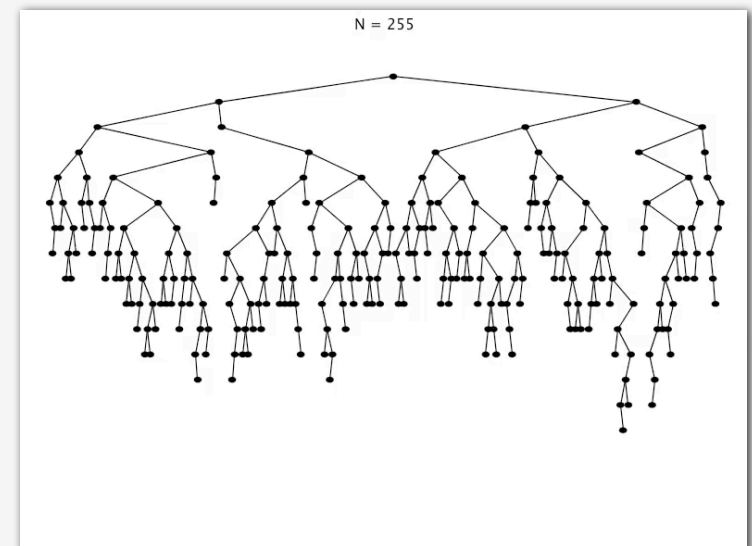
Observation. If keys inserted in random order, tree stays flat.



11

BST insertion: random order visualization

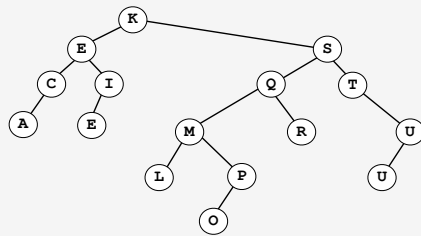
Ex. Insert keys in random order.



12

Correspondence between BSTs and quicksort partitioning

Q	U	I	C	K	S	O	R	T	E	X	A	M	P	L	E
E	R	A	T	E	S	L	P	U	I	M	Q	C	X	O	K
E	C	A	I	E	K	L	P	U	T	M	Q	R	X	O	S
A	C	E	I	E	K	L	P	U	T	M	Q	R	X	O	S
A	C	E	I	E	K	L	P	U	T	M	Q	R	X	O	S
A	C	E	I	E	K	L	P	U	T	M	Q	R	X	O	S
A	C	E	E	I	K	L	P	U	T	M	Q	R	X	O	S
A	C	E	E	I	K	L	P	O	R	M	Q	S	X	U	T
A	C	E	E	I	K	L	P	O	M	Q	R	S	X	U	T
A	C	E	E	I	K	L	M	O	P	Q	R	S	X	U	T
A	C	E	E	I	K	L	M	O	P	Q	R	S	X	U	T
A	C	E	E	I	K	L	M	O	P	Q	R	S	X	U	T
A	C	E	E	I	K	L	M	O	P	Q	R	S	X	U	T
A	C	E	E	I	K	L	M	O	P	Q	R	S	X	U	T
A	C	E	E	I	K	L	M	O	P	Q	R	S	T	U	X
A	C	E	E	I	K	L	M	O	P	Q	R	S	T	U	X
A	C	E	E	I	K	L	M	O	P	Q	R	S	X	U	T
A	C	E	E	I	K	L	M	O	P	Q	R	S	T	U	X



Remark. Correspondence is 1-1 if no duplicate keys.

13

BSTs: mathematical analysis

Proposition. If keys are inserted in **random** order, the expected number of compares for a search/insert is $\sim 2 \ln N$.

Pf. 1-1 correspondence with quicksort partitioning.

Proposition. [Reed, 2003] If keys are inserted in random order, expected height of tree is $\sim 4.311 \ln N$.

But... Worst-case for search/insert/height is N .
(exponentially small chance when keys are inserted in random order)

14

ST implementations: summary

implementation	guarantee		average case		ordered ops?	operations on keys
	search	insert	search hit	insert		
unordered array	N	N	$N/2$	N	no	<code>equals()</code>
unordered list	N	N	$N/2$	N	no	<code>equals()</code>
ordered array	$\lg N$	N	$\lg N$	$N/2$	yes	<code>compareTo()</code>
ordered list	N	N	$N/2$	$N/2$	yes	<code>compareTo()</code>
BST	N	N	$1.39 \lg N$	$1.39 \lg N$?	<code>compareTo()</code>

Next challenge. Ordered symbol tables ops.

15

- › basic implementations
- › randomized BSTs
- › ordered symbol table ops

16

Ordered symbol table operations

Minimum. Smallest key in table.

Maximum. Largest key in table.

Floor. Largest key \leq to a given key.

Ceiling. Smallest key \geq to a given key.

Rank. Number of keys $<$ than given key.

Select. Key of given rank.

Size. Number of keys in a given range.

Iterator. All keys in order.

	keys	values
<code>min()</code>	09:00:00	Chicago
	09:00:03	Phoenix
	09:00:13	Houston
<code>get(09:00:13)</code>	09:00:59	Chicago
	09:01:10	Houston
<code>floor(09:05:00)</code>	09:03:13	Chicago
	09:10:11	Seattle
<code>select(7)</code>	09:10:25	Seattle
	09:14:25	Phoenix
	09:19:32	Chicago
	09:19:46	Chicago
<code>keys(09:15:00, 09:25:00)</code>	09:21:05	Chicago
	09:22:43	Seattle
	09:22:54	Seattle
	09:25:52	Chicago
<code>ceiling(09:30:00)</code>	09:35:21	Chicago
	09:36:14	Seattle
<code>max()</code>	09:37:44	Phoenix

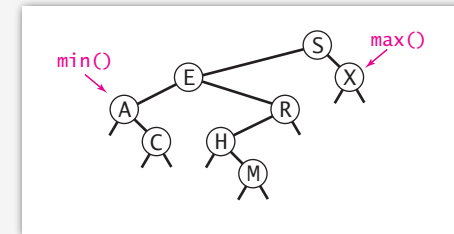
`size(09:15:00, 09:25:00) is 5`
`rank(09:10:25) is 7`

17

Minimum and maximum

Minimum. Smallest key in table.

Maximum. Largest key in table.



Q. How to find the min / max.

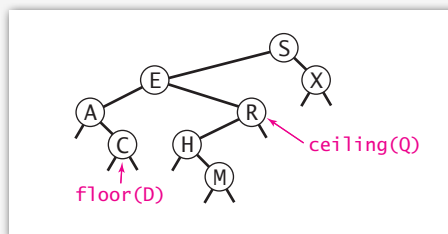
A.

18

Floor and ceiling

Floor. Largest key \leq to a given key.

Ceiling. Smallest key \geq to a given key.



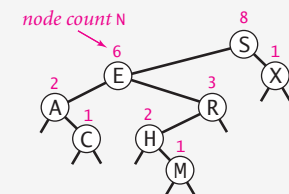
Q. How to find the floor /ceiling.

A.

19

Rank

Rank. How many keys $<$ k ?



```
public int rank(Key key)
{ return rank(key, root); }

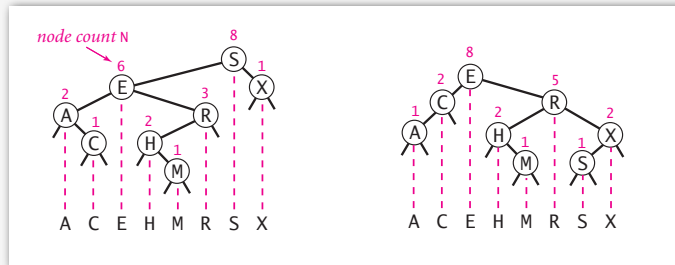
private int rank(Key key, Node x)
{
    if (x == null) return 0;
    int cmp = key.compareTo(x.key);
    if (cmp < 0) return rank(key, x.left);
    else if (cmp > 0) return 1 + size(x.left) + rank(key, x.right);
    else return size(x.left);
}
```

how to implement size() efficiently?

20

Subtree counts

In each node, we store the number of nodes in the subtree rooted at that node.



Remark. This facilitates efficient implementation of `rank()` and `select()`.

21

BST implementation: subtree counts

```
private class Node
{
    private Key key;
    private Value val;
    private Node left;
    private Node right;
    private int N;
}
```

```
public int size()
{ return size(root); }

private int size(Node x)
{
    if (x == null) return 0;
    return x.N;
}
```

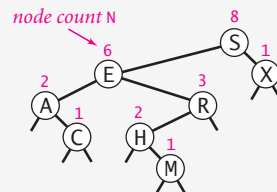
nodes in subtree

```
private Node put(Node x, Key key, Value val)
{
    if (x == null) return new Node(key, val);
    int cmp = key.compareTo(x.key);
    if (cmp < 0) x.left = put(x.left, key, val);
    else if (cmp > 0) x.right = put(x.right, key, val);
    else if (cmp == 0) x.val = val;
    x.N = 1 + size(x.left) + size(x.right);
    return x;
}
```

22

Range count

Range count. How many keys between `lo` and `hi`?



```
public int size(Key lo, Key hi)
{
    if (contains(hi)) return rank(hi) - rank(lo) - 1;
    else return rank(hi) - rank(lo);
}
```

number of keys < hi

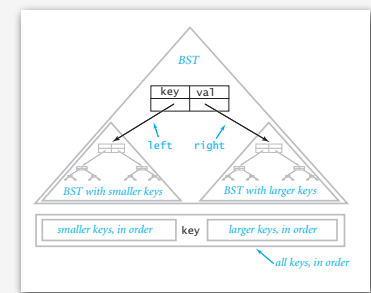
23

Inorder traversal

- Traverse left subtree.
- Enqueue key.
- Traverse right subtree.

```
public Iterable<Key> allKeys()
{
    Queue<Key> q = new Queue<Key>();
    inorder(root, queue);
    return q;
}

private void inorder(Node x, Queue<Key> q)
{
    if (x == null) return;
    inorder(x.left, q);
    q.enqueue(x.key);
    inorder(x.right, q);
}
```



Property. Inorder traversal of a BST yields keys in ascending order.

24

Inorder traversal

- Traverse left subtree.
- Enqueue key.
- Traverse right subtree.

```

visit(S)
  visit(E)
    visit(A)
      enqueue A
    visit(C)
      enqueue C
    enqueue E
  visit(R)
    visit(H)
      enqueue H
    visit(M)
      enqueue M
  print R
  enqueue S
visit(X)
  enqueue X
    
```

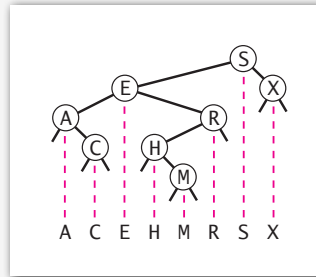
recursive calls

queue

function call stack

```

S
S E
S E A
S E A C
S E R
S E R H
S E R H M
R
S
S X
X
    
```



25

ST implementations: summary

implementation	guarantee			average case			ordered iteration?	operations on keys
	search	insert	delete	search hit	insert	delete		
sequential search (linked list)	N	N	N	N/2	N	N/2	no	equals ()
binary search (ordered array)	lg N	N	N	lg N	N/2	N/2	yes	compareTo ()
BST	N	N	N	1.39 lg N	1.39 lg N	?	yes	compareTo ()

Next lecture. Can we guaranteed performance?

26

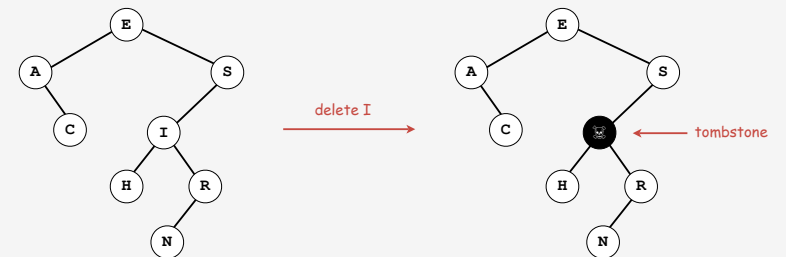
- basic implementations
- randomized BSTs
- deletion in BSTs

27

BST deletion: lazy approach

To remove a node with a given key:

- Set its value to `null`.
- Leave key in tree to guide searches (but don't consider it equal to search key).



Cost. $O(\log N')$ per insert, search, and delete (if keys in random order), where N' is the number of elements ever inserted in the BST.

Unsatisfactory solution. Tombstone overload.

28

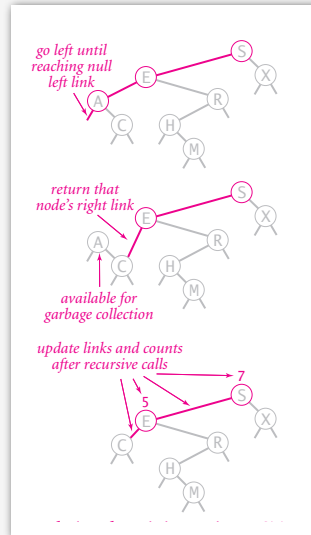
Deleting the minimum

To delete the minimum key:

- Go left until finding a node with a null left link.
- Replace that node by its right link.
- Update subtree counts.

```
public void deleteMin()
{ root = deleteMin(root); }

private Node deleteMin(Node x)
{
    if (x.left == null) return x.right;
    x.left = deleteMin(x.left);
    x.N = 1 + size(x.left) + size(x.right);
    return x;
}
```

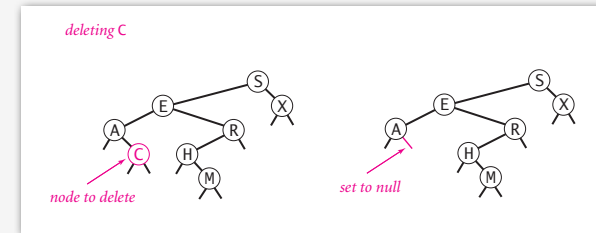


29

Hibbard deletion

To delete a node with key k: search for node t containing key k.

Case 0. [0 children] Delete t by setting parent link to null.

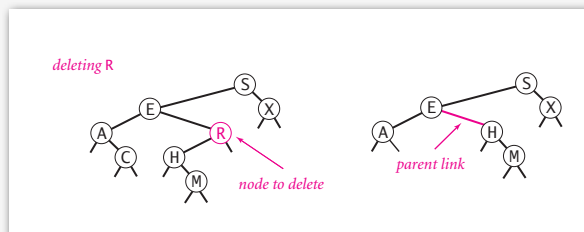


30

Hibbard deletion

To delete a node with key k: search for node t containing key k.

Case 1. [1 child] Delete t by replacing parent link.



31

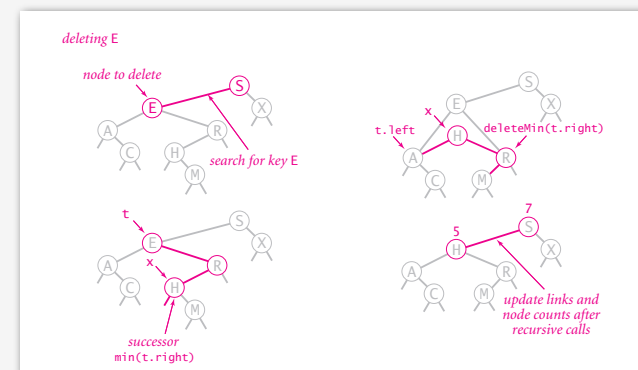
Hibbard deletion

To delete a node with key k: search for node t containing key k.

Case 2. [2 children]

- Find successor x of t.
- Delete the minimum in t's right subtree.
- Put x in t's spot.

- ← x has no left child
- ← but don't garbage collect x
- ← still a BST



32

Hibbard deletion: Java implementation

```

public void delete(Key key)
{ root = delete(root, key); }

private Node delete(Node x, Key key) {
    if (x == null) return null;
    int cmp = key.compareTo(x.key);
    if (cmp < 0) x.left = delete(x.left, key);
    else if (cmp > 0) x.right = delete(x.right, key);
    else {
        if (x.right == null) return x.left;

        Node t = x;
        x = min(t.right);
        x.right = deleteMin(t.right);
        x.left = t.left;
    }
    x.N = size(x.left) + size(x.right) + 1;
    return x;
}

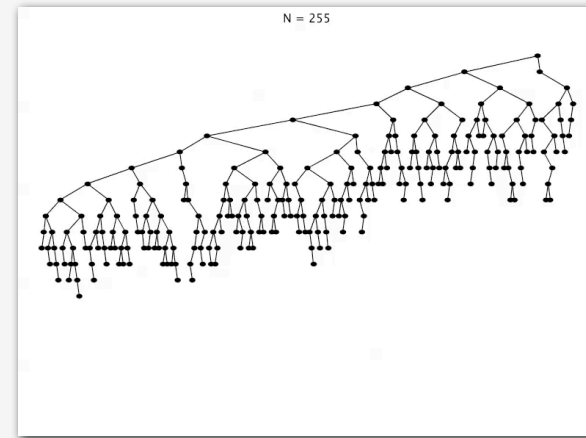
```

← search for key
← no right child
← replace with successor
← update subtree counts

33

Hibbard deletion: analysis

Unsatisfactory solution. Not symmetric.



Surprising consequence. Trees not random (!) \Rightarrow $\sqrt{\text{N}}$ per op.
 Longstanding open problem. Simple and efficient delete for BSTs.

34

ST implementations: summary

implementation	guarantee			average case			ordered iteration?	operations on keys
	search	insert	delete	search hit	insert	delete		
sequential search (linked list)	N	N	N	N/2	N	N/2	no	equals ()
binary search (ordered array)	lg N	N	N	lg N	N/2	N/2	yes	compareTo ()
BST	N	N	N	1.39 lg N	1.39 lg N	\sqrt{N}	yes	compareTo ()

↖ ↗
 other operations also become \sqrt{N}
 if deletions allowed

Next lecture. Guarantee logarithmic performance for all operations.

35