

ADVANCED DATA STRUCTURES

Lectures 1 and 2

Introduction. Binary search trees and red-black trees.
Efficiency issues

Mircea Marin

mircea.marin@e-uvf.ro

Organizatorial items

- Lecturer and Teaching Assistant: Mircea Marin
 - email: `mircea.marin@e-uvv.ro`
- Course objectives:
 - 1 Become familiar with some of the advanced data structures (ADS) and algorithms which underpin much of today's computer programming
 - 2 Recognize the data structure and algorithms that are best suited to model and solve your problem
 - 3 Be able to perform a qualitative analysis of your algorithms – time complexity analysis
- Course webpage:
<http://web.info.uvt.ro/~mmarin/lectures/ADS>
- Handouts: will be posted on the webpage of the lecture
- Grading: 40% final exam (written), 60% labwork (mini-projects)
- Attendance: required

Organizatorial items

- Lab work: implementation in C++ of applications, using data structures and algorithms presented in this lecture
- Requirements:
 - ▷ Be up-to-date with the presented material
 - ▷ Prepare the programming assignments for the stated deadlines
- Prerequisites:
 - 1 A good understanding of the data structures and related algorithms, taught in the lecture DATA STRUCTURES.
 - 2 Familiarity with the C++ programming language.
 - 3 Recommended IDEs: Eclipse or Code::Blocks
- Recommended textbook:
 - Cormen, Leiserson, Rivest. *Introduction to Algorithms*. MIT Press.

Dynamic sets

Elements of a dynamic set

A **dynamic set** is a collection of **objects** that may grow, shrink, or otherwise change over time.

- Objects have **fields**.
- A **key** is a field that uniquely identifies an object.
 - ▶ In all implementations of dynamic sets presented in this lecture, keys are assumed to be totally ordered.
- Operations on a dynamic set are grouped into two categories:
 - 1 **Queries:** they simply return information about the set.
Typical examples: **SEARCH**, **MINIMUM**, **MAXIMUM**, **SUCCESSOR**, **PREDECESSOR**
 - 2 **Modifying operations:** **INSERT**, **DELETE**

Operations on a dynamic set S

- 1 **search**(S, k): returns a pointer p to an element in S such that $p \rightarrow key = k$, or Nil if no such element exists.
- 2 **minimum**(S): returns a pointer to the element of S whose key is minimum, or Nil if S has no elements.
- 3 **maximum**(S): returns a pointer to the element of S whose key is maximum, or Nil if S has no elements.
- 4 **successor**(S, x): returns the next element larger than x in S , or Nil if x is the maximum element.
- 5 **predecessor**(S, x): returns the next element smaller than x in S , or Nil if x is the minimum element.
- 6 **insert**(S, p): augment S with the element pointed to by p .
If p is Nil, do nothing.
- 7 **del**(S, p): remove the element pointed to by p from S .
If p is Nil, do nothing.

Operations 1-5 are **queries**; operations 6-7 are **modifying operations**.

Dynamic sets

Implementation issues

- ▶ The complexity of operations of a dynamic set is measured in terms of its **size** $n =$ number of elements.
- ▶ Different implementations of dynamic sets vary by the runtime complexity of their operations
 - The choice of an implementation depends on the operations we perform most often.
 - Typical examples:
 - Binary search trees
 - Red-black trees
 - B-trees
 - Binomial heaps
 - Fibonacci heaps
 - ...

Dynamic sets for efficient search

1. Binary search trees (Recap)


- Relevant information is stored in the nodes of a **binary tree**
 - ▶ every non-leaf node has a left and a right child.
- Every node has a **unique key**, that is used to identify a node:
- The finding of a node with a particular key must be **fast**
⇒ the keys are distributed in a special way:
 - For every node, its key is larger than the keys of the nodes to its left, and smaller than the keys of the nodes to its right.

Binary search trees

C++ classes

```
struct Node {
    int key;           // key
    Node *p;          // pointer to parent
    Node *left;       // pointer to left child
    Node *right;      // pointer to right child
    ...               // satellite data
};

struct BSTree {
    Node *root;       // pointer to node at root position
    ...               // operations on trees
};
```




Leaf nodes are empty: they contain no data

Binary search trees

C++ classes

```
struct Node {  
    int key;           // key  
    Node *p;          // pointer to parent  
    Node *left;       // pointer to left child  
    Node *right;      // pointer to right child  
    ...               // satellite data  
};  
struct BSTree {  
    Node *root;       // pointer to node at root position  
    ...              // operations on trees  
};
```



Leaf nodes are empty: they contain no data


Requirement: $k_l \leq k \leq k_r$

Binary search trees

C++ classes

```
struct Node {
    int key;           // key
    Node *p;          // pointer to parent
    Node *left;       // pointer to left child
    Node *right;      // pointer to right child
    ...               // satellite data
};

struct BSTree {
    Node *root;       // pointer to node at root position
    ...               // operations on trees
};
```



Leaf nodes are empty: they contain no data

Requirement: $k_l \leq k \leq k_r$

If n is the root node then $n.p == 0$.


Binary search trees

Implementation issues and diagrammatic representations


In OOP (e.g., C++), there are two ways to represent leaf nodes:

- 1 with the null pointer
- 2 with a single sentinel node NIL, which is also the root's parent.

Example (Diagrammatic representations)



or



Binary search trees

Time complexity of operations

ASSUMPTION: S is an instance of `BSTree`, which represents a binary search tree with n nodes.

operation	average case	worst case
<code>S.search(k)</code>	$O(\log_2 n)$	$O(n)$
<code>S.minimum()</code>	$O(\log_2 n)$	$O(n)$
<code>S.maximum()</code>	$O(\log_2 n)$	$O(n)$
<code>S.successor(x)</code>	$O(\log_2 n)$	$O(n)$
<code>S.predecessor(x)</code>	$O(\log_2 n)$	$O(n)$
<code>S.insert(x)</code>	$O(\log_2 n)$	$O(n)$
<code>S.del(x)</code>	$O(\log_2 n)$	$O(n)$

Binary search trees

- Average case time complexity is $O(\log n)$ for all operations.
- Worst case time complexity is $O(n)$ for all operations.

Question: Can we improve the worst case time complexity of binary search trees without many changes?

Answer: Yes; **Red-black trees**

Binary Search trees versus Red-Black trees

Node structures

Binary search tree fields	Red-black tree fields
<code>key</code>	<code>key</code>
<code>left, right</code> : pointers to children	<code>left, right</code> : pointers to children
<code>p</code> : pointer to parent	<code>p</code> : pointer to parent
	<code>color</code> : RED or BLACK

Binary Search trees versus Red-Black trees

Tree structures

- Similarities: both kinds of trees satisfy the binary search tree property:
 - if x is a node with left child y then $y.key \leq x.key$
 - if x is a node with right child y then $x.key \leq y.key$
- Properties specific to red-black trees (the **red-black properties**)
 - 1 Every node is either red or black
 - 2 The root is black
 - 3 Every leaf NIL is black.
 - 4 The children of a red node are black.
 - 5 Every path from a node to a descendant leaf contains the same number of black nodes.

- For binary-search trees, no memory needs to be allocated for leaf nodes. All pointers to a leaf node are assumed to be `Nil`, which is 0.
- For Red-black trees, it is convenient to represent all leaves with a black sentinel node `NIL`:
 - The left child, right child, and parent of `NIL` are `NIL`.
 - Also, the parent of the root node of a red-black tree is assumed to be `NIL`.

Red-black trees

Important notions and derived properties

NOTIONS:

Height $h(n)$ of a node n = number of edges in a longest path from n to a leaf.


Black height $bh(n)$ of a node n = number of black nodes (including NIL) on the path from n to a leaf, not counting n .

DERIVED PROPERTIES:


- 1 $bh(n) \geq h(n)/2$ for every node n .
- 2 The subtree rooted at any node x contains $\geq 2^{bh(x)} - 1$ internal nodes.
- 3 A red-black tree with n internal nodes has height $\leq 2 \log_2(n + 1)$.

Binary search trees and Red-black trees

Graphical representation



(a) Binary search tree



(b) Red-Black tree

- Nodes are represented by circles with the key value written inside.
 - Nodes may contain more data, not shown in the graphical representation.
 - For Red-Black trees, the circles are colored with the node colour.
- Thick lines between nodes are the pointers from parent to children.
- Dashed arrows are pointers from node to parent. They are usually not drawn.

Tree representations of a dynamic set S

Encodings in C++

As a binary search tree:

```
class Node {
public:
    int key;
    Node* left;
    Node* right;
    Node* p;
    // constructors
    ...
}

class BSTree {
public:
    Node* root; // pointer to root
    BSTree() { root = 0; }
    ...
}
```

As a red-black tree:

```
class RBNode {
public:
    int key;
    RBNode* left;
    RBNode* right;
    RBNode* p;
    enum Color {RED, BLACK};
    Color color;
    // constructors
    ...
}

class RBTree {
public:
    RBNode* root; // pointer to root
    RBTree() { root = 0; }
    ...
}
```

Operations on Red-black trees

ASSUMPTION: S is an instance of `RBTtree`, which represents a red-black search tree with n nodes.

operation	average case	worst case
$S.search(k)$	$O(\log_2 n)$	$O(\log_2 n)$
$S.minimum()$	$O(\log_2 n)$	$O(\log_2 n)$
$S.maximum()$	$O(\log_2 n)$	$O(\log_2 n)$
$S.successor(x)$	$O(\log_2 n)$	$O(\log_2 n)$
$S.predecessor(x)$	$O(\log_2 n)$	$O(\log_2 n)$

- These are query operations, implemented the same way as for binary search trees.
- The modifying operations `insert` and `del` must be redesigned carefully, to guarantee that the newly produced tree has the red-black properties.

Insertion in binary search trees

`T.insert(z)` where

`T` : binary search tree

`z` : pointer to node with `z->key = v`,

`z->left = z->right = Nil`.

Effect: `T` and `z` are modified such that `z` is inserted at the right position in `T`.

```
insert(Node* z) // method of class BSTree
```

```
1 y = Nil
2 x = root
3 while x ≠ Nil
4   y = x
5   if z->key < x->key
6     x = x->left
7   else x = x->right
8 z->p = y
9 if y == Nil
10  root = z
11 else if (z->key < y->key)
12  y->left = z
13 else y->right = z
```

Deletion from binary search trees


`T.del(z)` deletes node `z` from the binary search tree `T`

```
del(Node* z) // method of class BSTree
1 if z->left == 0 or z->right == 0
2   y := z
3 else y = successor(z)
4 if y->left != 0
5   x = y->left
6 else x = y->right
7 if x != 0
8   x->p = y->p
9 if y->p == 0
10  root = x
11 else if y == y->p->left
12     y->p->left = x
13     else y->p->right = x
14 if y != z
15     z->key = y->key
16     if y has other fields, copy them to z too
17 return y
```

Binary search trees

Case 1: Deletion of a node z without children

Example




- Node with key 13 has no children
⇒ we simply remove it from the binary search tree.

Binary search trees

Case 2: Deletion of a node z with one child

Example




Node z with key 16 has only one child



\Rightarrow the child of z becomes the child of the parent of z

Binary search trees

Case 2: Deletion of a node z with two children



$\text{del}(z) \rightarrow$



Red-Black trees

Example





Figure: A balanced RB-tree with black-height 3

Insertion and deletion in red-black trees

Remarks

- The operations `insert` and `del` for binary search trees can be run on red-black trees with n keys
 - ⇒ they take $O(\log_2(n))$ time.
 - ⇒ they may destroy the red-black properties of the tree:
 - `insert` of RED node might violate property 4; of BLACK node might violate property 5.
 - `del` of RED node: no property violations; of BLACK node might violate properties 2, 4, 5.
- ⇒ **the red-black properties must be restored:**
- ▶ Some nodes must change color
 - ▶ Some pointers must be changed

LeftRotate and RightRotate




The rotation operations on a binary search tree.

$T.RightRotate(x)$ transforms the configuration of the two nodes on the left into the configuration on the right by changing a constant number of pointers.

The configuration on the right can be transformed into the configuration on the left by the inverse operation $T.LeftRotate(y)$. The two nodes may appear anywhere in a binary search tree T . The letters α , β , and γ represent binary subtrees.

LeftRotate and RightRotate



The rotation operations on a binary search tree.

$T.RightRotate(x)$ transforms the configuration of the two nodes on the left into the configuration on the right by changing a constant number of pointers.

The configuration on the right can be transformed into the configuration on the left by the inverse operation $T.LeftRotate(y)$. The two nodes may appear anywhere in a binary search tree T . The letters α , β , and γ represent binary subtrees.

A rotation operation preserves the inorder ordering of keys: the keys in α precede $x \rightarrow key$, which precedes the keys in β , which precede $y \rightarrow key$, which precedes the keys in γ .

LeftRotate and RightRotate

Assumption: $x \rightarrow \text{right} \neq \text{NIL}$.

```
LeftRotate(Node* x) // method of class BSTree
```

```
1 y = x->right
2 x->right = y->left
3 if y->left  $\neq$  Nil
4   y->left->p = x
5 y->p = x->p
6 if x->p == Nil
7   root = y
8 else if x == x->p->left
9   x->p->left = y
10 else x->p->right = y
11 y->left = x
12 x->p = y
```

- The code for `RightRotate` is similar.
- Both `LeftRotate` and `RightRotate` run in $O(1)$ time.

LEFTROTATE illustrated

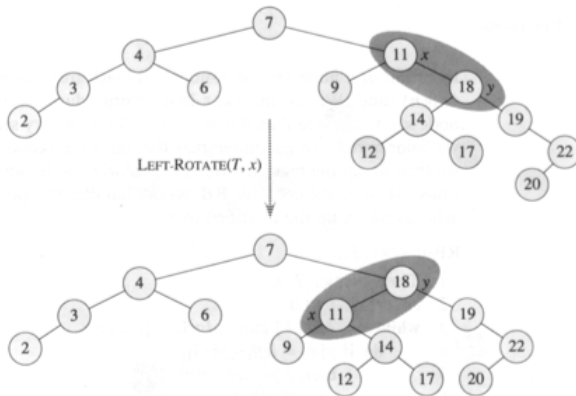


Figure 14.3 An example of how the procedure $\text{LEFT-ROTATE}(T, x)$ modifies a binary search tree. The NIL leaves are omitted. Inorder tree walks of the input tree and the modified tree produce the same listing of key values.

Insertion

Main idea

To insert a node x in a red-black tree T , proceed as follows:

- Perform insertion of x in T , as if T were a binary search tree.
- Color x to be **RED**.
- Adjust the color of the modified tree, by recoloring nodes and performing rotations.
- These ideas are implemented in the `RBInsert` procedure.

RB-trees

Insertion

```
RBInsert(Node* z) // method of class RBNode
1 y = NIL
2 x = root
3 while x ≠ Nil
4     y = x
5     if z->key < x->key
6         then x = x->left
7         else x = x->right
8 z->p = y
9 if y == NIL
10    then root = z
11    else if z->key < y->key
12        then y->left = z
13        else y->right = z
14 z->left = NIL
15 z->right = NIL
16 z->color = RED
17 RBInsertFixup(z)
```

Analysis of RBInsert

- `RBInsert` ends by coloring the new node z red.
- Then it calls `RBInsertFixup` because we could have violated some red-black properties:
 - The red-black properties 1, 3, and 5 are not violated.
 - Property 2 is violated if z is the root.
 - Property 4 is violated if $z \rightarrow p$ is **RED**, because both z and $z \rightarrow p$ are **RED**.

Analysis of `RBInsert`

- `RBInsert` ends by coloring the new node z red.
- Then it calls `RBInsertFixup` because we could have violated some red-black properties:
 - The red-black properties 1, 3, and 5 are not violated.
 - Property 2 is violated if z is the root.
 - Property 4 is violated if $z \rightarrow p$ is **RED**, because both z and $z \rightarrow p$ are **RED**.

These violations are removed by calling `RBInsertFixup`.

RBInsertFixup

```
RBInsertFixup(Node* z) // method of class RBNode
1 while z->p->color == RED
2     if z->p == z->p->p->left
3     then y = z->p->p->right
4         if y->color == RED
5             then z->p->color = BLACK // Case 1
6                 y->color = BLACK // Case 1
7                 z->p->p->color = RED // Case 1
8                 z = z->p->p // Case 1
9         else if z == z->p->right
10            then z = z->p // Case 2
11                LeftRotate(z) // Case 2
12                z->p->color = BLACK // Case 3
13                z->p->p->color = RED // Case 3
14                RightRotate(z->p->p) // Case 3
15     else
16         (same as then clause with right and left exchanged)
17 root->color = BLACK
```

RBInsertFixup

```
RBInsertFixup(Node* z) // method of class RBNode
1 while z->p->color == RED
2     if z->p == z->p->p->left
3     then y = z->p->p->right
4         if y->color == RED
5             then z->p->color = BLACK           // Case 1
6                 y->color = BLACK             // Case 1
7                 z->p->p->color = RED          // Case 1
8                 z = z->p->p                  // Case 1
9             else if z == z->p->right
10                then z = z->p                // Case 2
11                    LeftRotate(z)          // Case 2
12                    z->p->color = BLACK      // Case 3
13                    z->p->p->color = RED      // Case 3
14                    RightRotate(z->p->p)    // Case 3
15        else
16            (same as then clause with right and left exchanged)
17 root->color = BLACK
```

Remark: The following loop invariant holds at the start of each **while** loop:


- 1 z is RED
- 2 There is at most one red-black violation: z is RED (property 2), or z and z->p are both RED

RBInsertFixup(Z)


Diagrammatic description

The new node z is **RED**, and inserted at the bottom of tree T . If the parent of z is red, we must fix T .

- Case 1: parent and uncle of z are **RED**




- Case 2: z is right child, and its uncle is **BLACK**



RBInsertFixup (Z)

Diagrammatic description (continued)

- Case 3: z is left child, and its uncle is BLACK



Deletion

- We will present a procedure $T.RBDelete(x)$ that performs deletion of node x from RB-tree T in $O(\log_2(n))$ time.
- $T.RBDelete(x)$ is a subtle adjustment of the deletion procedure for binary search trees. After splicing out a node, it calls an auxiliary procedure $RBDeleteFixup(x)$ that changes colors and performs rotations to restore the red-black properties.

RBDelete versus del

RBDelete(z)

```
1 if z->left == Nil or z->right == Nil
2  y := z
3 else y = successor(z)
4 if y->left ≠ Nil
5  x = y->left
6 else x = y->right
7 x->p = y->p
8 if y->p = Nil
9  root = x
10 else if y == y->p->left
11  y->p->left = x
12 else y->p->right = x
13 if y ≠ z
14  z->key = y->key
15  if y has other fields,
    copy them to z too
16 if y->color = BLACK
17  RBDeleteFixup(x)
18 return y
```

DEL(z)

```
1 if z->left == Nil or z->right == Nil
2  y := z
3 else y = successor(z)
4 if y->left ≠ Nil
5. x = y->left
6 else x = y->right
7 if x ≠ Nil
8  x->p = y->p
9 if y->p = Nil
10 root = x
11 else if y == y->p->left
12  y->p->left = x
13 else y->p->right = x
14 if y ≠ z
15  z->key = y->key
16  if y has other fields,
    copy them to z too
17 return y
```

$T.RBDelete(z)$ returns node y which is spliced out from T .

- ▷ If y is red, the tree without y is red-black.
- ▷ Otherwise, y is black and the tree without y may violate the following red-black properties:
 - property 2: when y is the root, and y has only one child, which is red.
 - property 4: when x and $y \rightarrow p$ (which becomes $x \rightarrow p$) are both red.
 - property 5: all paths that previously contained y have one fewer black nodes.

To restore the red-black properties, we call `RBDeleteFixup(x)` where x is either

- the sole child of y , before y was spliced out, or
- the sentinel node *Nil*

Note: In both cases, $x \rightarrow p = y \rightarrow p$

RBDELETEFIXUP(x)


```
RBDeleteFixup(x)
1 while x ≠ root and x->color == BLACK
2   if x == x->p->left
3     w = x->p->right
4     if w->color = RED
5       w->color = BLACK           Case 1
6       x->p->color = RED           Case 1
7       LeftRotate(x->p)          Case 1
8       w = x->p->right            Case 1
9     if w->left->color == BLACK and w->right->color == BLACK
10      w->color = RED             Case 2
11      x = x->p                   Case 2
12    else if w->right->color == BLACK
13      w->left->color = BLACK      Case 3
14      w->color = RED             Case 3
15      RightRotate(w)            Case 3
16      w = x->p->right            Case 3
17      w->color = x->p->color       Case 4
18      x->p->color = BLACK          Case 4
19      w->right->color = BLACK     Case 4
20      LeftRotate(x->p)           Case 4
21      x = root                  Case 4
22    else (same as then clause, with right and left exchanged)
23 x->color = BLACK
```

- ▷ Lines 1-22 are intended to move the extra black up the tree until either
 - 1 x points to a red node \Rightarrow we will color the node black (line 23)
 - 2 x points to the root \Rightarrow the extra black can be simply "removed"
 - 3 Suitable rotations and recolorings can be performed.

- ▷ Lines 1-22 are intended to move the extra black up the tree until either
 - ① x points to a red node \Rightarrow we will color the node black (line 23)
 - ② x points to the root \Rightarrow the extra black can be simply "removed"
 - ③ Suitable rotations and recolorings can be performed.
- ▷ In lines 1-22, x always points to a non-root black node that has the extra black, and w is set to point to the sibling of x . The **while** loop distinguishes 4 cases. See figure on next slide, where:
 - *Darkened nodes* are black, *heavily shaded nodes* are red, and *lightly shaded nodes* can be either red or black.
 - Small Greek letters represent arbitrary subtrees.
 - A node pointed to by x has an extra black.

RBDeleteFixup (X)

Implementation analysis



The four cases in the **while** loop:

- The only case that can cause the loop to repeat is case 2.
- (a) Case 1 is transformed into case 2,3, or 4 by exchanging colors of nodes B and D and performing a left rotation.
- (b) In case 2, the extra black represented by the pointer x is moved up the tree by coloring node D red and setting x to point to B . If we enter case 2 through case 1, the **while** loop terminates, since the color c is red.
- (c) Case 3 is transformed to case 4 by exchanging the colors of nodes C and D and performing a right rotation.
- (d) In case 4, the extra black represented by x can be removed by changing some colors and performing a left rotation.

- The height of the RB-tree is $O(\log_2(n)) \Rightarrow$ the total cost of the procedure without the call to `RBDeleteFixup` is $O(\log_2(n))$.
 - Within the `RBDeleteFixup` call, cases 1, 3, 4 each terminate after a constant number of color changes and ≤ 3 rotations. Case 2 is the only case in which the **while** loop can be repeated, and then the pointer x moves upward at most $O(\log_2(n))$ times, and no rotations are performed. Thus
 - \Rightarrow `RBDeleteFixup(x)` takes $O(\log_2(n))$ time and performs at most 3 rotations.
- \Rightarrow The overall time of `RBDelete(x)` is also $O(\log_2(n))$.

- Chapters 13 (Binary Search Trees), 14 (Red-Black trees), and Section 5.5 from the book
 - Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest. *Introduction to Algorithms*. McGraw Hill, 2000.