

Exam

Data structures DIT960

Time	Thursday 30th May 2013, 14:00–18:00
Place	V-huset
Course responsible	Nick Smallbone, tel. 0707 183062

The exam consists of **six easy questions (nos. 1–6)**, **two hard questions (nos. 7–8)** and **one extra hard question (no. 9)**, which counts as two hard questions.

For *Godkänd* you need to answer at least **four easy questions** correctly.

For *Väl Godkänd* you *also* need to answer **two hard questions** correctly.

Answering question 9 correctly counts as two hard questions; a partial answer to question 9 may count as one hard question.

Allowed aids One A4 piece of paper of notes, which may be hand-written or typed. You may write on both sides.

You may also bring a dictionary.

Note Begin each question on a new page.

Write your anonymous code (*not* your name) on every page.

Excessively complicated answers might be rejected.

Write legibly!

1. Suppose you have the following hash table, implemented using *linear probing*. For the hash function we are using the identity function, $h(x) = x$.

0	1	2	3	4	5	6	7	8
9	18		12	3	14	4	21	

- a) In which order could the elements have been added to the hash table? There are several correct answers, and you should choose them all.

- 9, 14, 4, 18, 12, 3, 21
- 12, 3, 14, 18, 4, 9, 21
- 12, 14, 3, 9, 4, 18, 21
- 9, 12, 14, 3, 4, 21, 18
- 12, 9, 18, 3, 14, 21, 4

Answer: 12, 14, 3, 9, 4, 18, 21, or 9, 12, 14, 3, 4, 21, 18.

- b) Add 30 to the hash table (assume there is no rehashing). What does the hash table look like now?

0	1	2	3	4	5	6	7	8
9	18		12	3	14	4	21	30

- c) Remove 3 from the hash table. What does it look like now?

0	1	2	3	4	5	6	7	8
9	18		12	XXX	14	4	21	30

Answer: 3 is replaced by a 'deleted' marker.

2. Implement binary search in Java. Your method should have the following type:

```
int search(Object[] array, Object key)
```

It should return the *index* of the object, or -1 if the object could not be found. It should take $O(\log n)$ time, and must not use recursion, or any Java standard library methods.

Warning: there are many different ways to write a *wrong* binary search! To pass this question, your code must be correct. *Make extra sure that you cannot get into an infinite loop.* You do not, however, have to consider integer overflow.

Answer on a separate sheet of paper. You may want to start from the following skeleton program:

```
int search(Object[] array, Object key) {
    int low = 0;
    int high = array.length - 1;
    while(...) {
        int mid = ...;
        if (array[mid].compare(key) < 0) {
            // array[mid] < key
            ...
        } else if (array[mid].compare(key) > 0) {
            // array[mid] > key
            ...
        } else {
            // array[mid] == key
            ...
        }
    }
    ...
}
```

One possible answer:


```
int search(Object[] array, Object key) {
    int low = 0;
    int high = array.length - 1;
    while(low <= high) {
        int mid = (low+high)/2;
        if (array[mid].compare(key) < 0) {
            // array[mid] < key
            low = mid+1;
        } else if (array[mid].compare(key) > 0) {
            // array[mid] > key
            high = mid-1;
        } else {
            // array[mid] == key
            return mid;
        }
    }
    return -1;
}
```

3. Which array out of A, B and C represents a binary heap? Only one answer is right.

	0	1	2	3	4	5	6	7	8	9	10	11
<input type="checkbox"/> A =	1	4	54	8	45	76	65	44	11	47	57	
<input type="checkbox"/> B =	1	8	45	13	43	47	44	65	23	76	57	
<input type="checkbox"/> C =	1	13	47	23	65	54	67	45	32	76	57	

Answer: A.

a) Write the heap out as a binary tree.



b) Add 46 to the heap, making sure to restore the heap invariant. How does the array look now?

	0	1	2	3	4	5	6	7	8	9	10	11
	1	4	46	8	45	54	65	44	11	47	57	76

4. Perform a *selection sort* on the following array. Your sort must be in-place, so you must not use a temporary array.

0	1	2	3	4	5	6	7	8
53	71	32	67	90	23	44	88	42

Show how the array looks after each execution of the outer loop. Also show which part of the array is sorted, for example by underlining it.

0	1	2	3	4	5	6	7	8
<u>23</u>	71	32	67	90	53	44	88	42

0	1	2	3	4	5	6	7	8
<u>23</u>	<u>32</u>	71	67	90	53	44	88	42

0	1	2	3	4	5	6	7	8
<u>23</u>	<u>32</u>	<u>42</u>	67	90	53	44	88	71

0	1	2	3	4	5	6	7	8
<u>23</u>	<u>32</u>	<u>42</u>	<u>44</u>	90	53	67	88	71

0	1	2	3	4	5	6	7	8
<u>23</u>	<u>32</u>	<u>42</u>	<u>44</u>	<u>53</u>	90	67	88	71

0	1	2	3	4	5	6	7	8
<u>23</u>	<u>32</u>	<u>42</u>	<u>44</u>	<u>53</u>	<u>67</u>	90	88	71


0	1	2	3	4	5	6	7	8
<u>23</u>	<u>32</u>	<u>42</u>	<u>44</u>	<u>53</u>	<u>67</u>	<u>71</u>	88	90

0	1	2	3	4	5	6	7	8
<u>23</u>	<u>32</u>	<u>42</u>	<u>44</u>	<u>53</u>	<u>67</u>	<u>71</u>	<u>88</u>	90

0	1	2	3	4	5	6	7	8
<u>23</u>	<u>32</u>	<u>42</u>	<u>44</u>	<u>53</u>	<u>67</u>	<u>71</u>	<u>88</u>	<u>90</u>


5. You are given the following AVL tree.

a) Mark each node with its AVL balance (left height minus right height).




b) Insert 60 into the tree, and balance it to restore the invariant. Write down the final tree.

60 becomes the left child of 62, and the root node now has a balance of -2. It is a right-left tree, and a pair of rotations fixes the invariant.



6. You are given the following weighted graph:



a) Perform Dijkstra's algorithm starting from node A. At each step the algorithm visits a new node. In which order are the nodes visited, and what is the distance to each of them? There are several possible orders to visit the nodes in – you may choose any of them.


Node	A	D	C	B	F	G	E	H
Distance	0	1	2	3	3	6	8	9

You can visit B and F in either order.

b) What is the shortest path to node **H**? List all the nodes along the way.

Distance 9 – A, C, F, G, E, H.

c) Use Prim's algorithm to construct a minimum spanning tree for the graph, starting from whichever node you like. Draw the tree below.



7. (**hard**) Your job is to implement a queue using a circular array. Your queue should implement the interface `Queue<E>` shown below, which is based on the Java queue interface.

```
public interface Queue<E> {  
    // Inserts the specified element at the back of the queue.  
    void add(E e);  
  
    // Retrieves and removes an element from the front of the  
    // queue, or returns null if the queue is empty.  
    E poll();  
  
    // Returns the number of elements in the queue.  
    int size();  
}
```

Note that:

- Your queue should not have any capacity restrictions.
- You should include a constructor for your queue.
- You may not use Java standard library functions in your solution.

Hint: you may find it useful to define a separate method `void resize()` that doubles the size of the array.

Answer on a separate sheet of paper.

One possible answer:

```
class CircularQueue<E> implements Queue<E> {
    E[] array = (E[]) new Object[1];
    int size = 0;
    int low = 0;
    int high = 0;

    void add(E e) {
        if (size == array.length) resize();
        high = (high+1) % array.length;
        size = size+1;
        array[high] = e;
    }

    E poll() {
        if (size == 0) return null;
        E result = array[low];
        low = (low+1) % array.length;
        size = size-1;
        return result;
    }

    void resize() {
        E[] newArray = (E[]) new Object(array.length * 2);
        for (int i = 0; i < size; i++) {
            newArray[i] = poll();
        }
        array = newArray;
        low = 0;
        high = size-1;
    }

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

8. (**hard**) Suppose we are given the following type of binary search trees in Haskell:

```
data Tree a = Nil | Node a (Tree a) (Tree a)
```

- a) Implement a function

```
greatest :: Ord a => Tree a -> a
```

that returns the greatest element in a non-empty binary search tree (in an empty binary search tree it is allowed to crash).

The complexity of your function should be $O(\text{height of tree})$, i.e., $O(\log n)$ for balanced trees, $O(n)$ for unbalanced trees.

- b) Write a function to delete an element. It will take two parameters, which are the element to delete and the tree, and have the following type:

```
delete :: Ord a => a -> Tree a -> Tree a
```

The complexity of your function should be $O(\text{height of tree})$, i.e., $O(\log n)$ for balanced trees, $O(n)$ for unbalanced trees.

Hint: it will help to use `greatest` when implementing `delete`.


Answer on a separate sheet of paper.

```
greatest :: Ord a => Tree a -> a
greatest (Node x _ Nil) = x
greatest (Node _ _ r) = greatest r

delete :: Ord a => a -> Tree a -> Tree a
delete x Nil = Nil
delete x (Node y l r)
  | x < y = Node y (delete x l) r
  | x > y = Node y l (delete x r)
-- otherwise, x == y
delete x (Node _ Nil r) = r
delete x (Node _ l r) =
  Node (greatest l) (delete (greatest l) l) r
```

9. (extra hard, worth two hard questions)

Suppose we want to add an array-like indexing operation to a binary tree. The idea is that `get(i)` will return the *i*th element that would be traversed in an in-order traversal of the tree. For example, in the tree to the right, `get(0)` would return `h`, `get(1)` would return `e`, `get(2)` would return `l`, `get(3)` would return `l`, `get(4)` would return `o`, and so on.



Your job is to efficiently implement `get`. Your implementation should have complexity $O(\text{height of tree})$, i.e. $O(\log n)$ for balanced trees, $O(n)$ for unbalanced trees. You will have to work out by yourself how to implement it. If you get part of the way, write down what you've found: you may get partial credit for it.

Here are some hints to help you:

- `get` should be a recursive function. It will either recurse into the left or right child, or return the value of the current node. First try to work out the *base case*: when is the current node the one we are looking for?
- To implement `get` efficiently, you will need to know the *size of each node*. Can you think how it helps to know the size of each node? Think especially about the size of the left child.
- If you are stuck, take the tree above, label each node with its size and index and look for a pattern – something that tells you which child to recurse into and with what index. Look at what happens with `get(3)`, `get(4)` and `get(5)`.
- Try writing down pseudocode once you have the basic idea.

You may choose either Java or Haskell. For Java, you should start from the following tree datatype (which records the size of each node) and skeleton implementation:

```
class Tree<E> {
  E value;
  int size; // The size of the node
  Tree<E> left, right;

  E get(int index) { /* your code goes here */ }
}
```

For Haskell, you should use the following datatype (which records the size of each node) and skeleton implementation:

```
data Tree a = Nil | Node Int a (Tree a) (Tree a)
  -- The 'Int' field is the size of the node
get :: Int -> Tree a -> a
get index tree = -- your code goes here
```

You may find it useful to write a function `size :: Tree a -> Int` that returns the size of any tree.

Answer on a separate sheet of paper.

The idea is that the size of the left subtree tells us where to find the index we are looking for:

- If the index is less than the size of the left subtree, the value we are looking for is inside the left subtree (and we should recurse into the left child)
- If the index is equal to the size of the left subtree, the current node is the one we are looking for
- If the index is greater than the size of the left subtree, the value we are looking for is inside the right subtree

Java implementation:

```
E get(int index) {
    int leftSize;
    if (left == null) leftSize = 0;
    else leftSize = left.size;


    if (index < leftSize)
        return left.get(index);
    else if (index == leftSize)
        return value;
    else
        return right.get(index - leftSize - 1);
}
```

Haskell implementation:

```
size :: Tree a -> Int
size Nil = 0
size (Node n _ _ _) = n

get :: Int -> Tree a -> a
get index (Node _ x l r)
    | index < size l = get index l
```

```
| index == size l = x  
| otherwise = get (index - size l - 1) r
```



<http://xkcd.com/835/>