

Midterm II Exam

15-122 Principles of Imperative Computation, Spring 2012
André Platzer Ananda Gunawardena

April 5, 2012

Name:

Andrew ID:

Instructions


- This exam is closed-book with one sheet of notes permitted.
- You have 80 minutes to complete the exam.
- There are 3 problems.
- Read each problem carefully before attempting to solve it.
- Do not spend too much time on any one problem.
- Consider if you might want to skip a problem on a first pass and return to it later.
- Consider writing out programs or diagrams on scratch paper first.
- And most importantly,

DON'T PANIC!


	Message Buffer	Virus Scan	BSTs	
	Prob 1	Prob 2	Prob 3	Total
Score				
Max	45	25	30	100
Grader				

1 Message Buffer (45 points)

You are implementing a network driver that stores the messages that it received in a queue. We can implement that queue in an array, remembering its length and two indices, one to the *front* of the queue and one to the *back*:



As usual, messages are inserted at the back and removed from the front. If either index ever passes the limit, it wraps around to the beginning of the array. You may assume that the limit is greater than zero. Due to this wrapping nature of the indices, the array can be visualized as a ring, which is why it is called a *ring buffer*. The *front* and *back* indices should only be equal if the queue is empty:



```
typedef struct queue* queue;
struct queue {
    int limit;
    message[] A;
    int front;
    int back;
};
```

As an example, here is the function that checks whether a queue is empty.

```
bool queue_empty(queue Q)
/*@requires is_queue(Q);
{
    return Q->front == Q->back;
}
```

(Continued)

Task 1 (10 pts). Identify all relevant data structure invariants for the (ring buffer) queue by implementing the `is_queue` function to check them.

```
bool is_queue(queue Q) {
```

```
}
```

Task 2 (5 pts). Implement the `deq` function that dequeues the message at the front of the queue.

```
message deq(queue Q)  
//@requires is_queue(Q);  
//@requires !queue_empty(Q);  
//@ensures is_queue(Q);  
{
```

```
}
```

Task 3 (5 pts). How can you use the idea behind how unbounded arrays double arrays to make sure the queue is never full, i.e., that `enq` always works successfully (without losing messages), no matter how big the array of the queue used to be before calling `enq`? You do not need to give code.

Task 4 (5 pts). Implement the `enq` function that enqueues a message. Make sure it works as expected no matter how big the array of the queue used to be. You can use helper functions that you have explained in your answer to Task 3.

```
void enq(queue Q, message e)
//@requires is_queue(Q);
//@ensures is_queue(Q);
//@ensures !queue_empty(Q);
{
```

```
}
```

(Continued)

The following questions refer to the `enq` function that you just implemented. Express your answer in big- O notation in terms of the number of array reads and writes that an operation performs.

Task 5 (5 pts). What is the worst-case asymptotic complexity of a single queue operation on a queue of size n ? Justify your answer in a sentence or two.

Task 6 (5 pts). What is the worst-case asymptotic complexity of a sequence of k queue operations starting from an empty queue? Justify your answer in a sentence or two.

(Continued)

Task 7 (10 pts). How can you change your implementation so that it gives priority to high-priority messages (flagged with an extra input `important==true` on delivery). The order of the messages upon delivery via `deq` should still be in order, except that high-priority messages (`important==true`) should always be delivered faster than low-priority messages (`important==false`). No code required. A *precise* explanation of your approach is sufficient. If you are unsure about whether your explanation is precise, you should write code.

```
void enqp(queue Q, message e, bool important)
  //@requires is_queue(Q);
  //@ensures is_queue(Q);
  //@ensures !queue_empty(Q);
  ;
```

2 Virus Scan (25 points)

You have a database of m bit sequences (called *signatures*), each of length at most k . Each signature is a bit sequence that identifies a computer virus. Your task is to write a function `infected` that checks whether one of the virus signatures is contained in a file of size n bits. We assume that the file has already been read into an array F of bits (represented as `bool`) of length n .

Task 1 (5 pts). A naive implementation of `infected` looks at each position i of F and checks whether any of the signatures start at that position by comparing the signature with F starting at i . What is the asymptotic number of comparisons of bits (from F), as a function of m, n and k , of this naive implementation of the function `infected`? Justify your answer (1 or 2 sentences).

(Continued)

Task 2 (15 pts). In order to improve the efficiency of your virus scanner, you decide to store the virus signatures in a binary trie. Recall that binary tries decompose bit sequence representations bit by bit in order to access data stored in the trie.

```
typedef struct btrie* btrie;
struct btrie {
    bool virus; /* true if identifies a virus */
    btrie btrue; /* trie for bit true */
    btrie bfalse; /* trie for bit false */
};
```

Reimplement `infected` in a more efficient way by using the virus signatures stored in a `btrie`. Your implementation does not have to achieve the best possible performance. It is enough if the code is correct and asymptotic complexity strictly better than that of the naive implementation.

Hint: If you use recursion, then you may find it useful to implement an auxiliary function that handles one index.

```
bool infected(btrie signatures, bit[] F, int i, int n)
//@requires 0 <= i && i <= n && n == \length(F);
{

}

}
```

Task 3 (5 pts). What is the asymptotic number of comparisons of bits (from F), as a function of m, n and k , of your improved implementation of the function `infected` from Task 2? Justify your answer in a sentence or two.

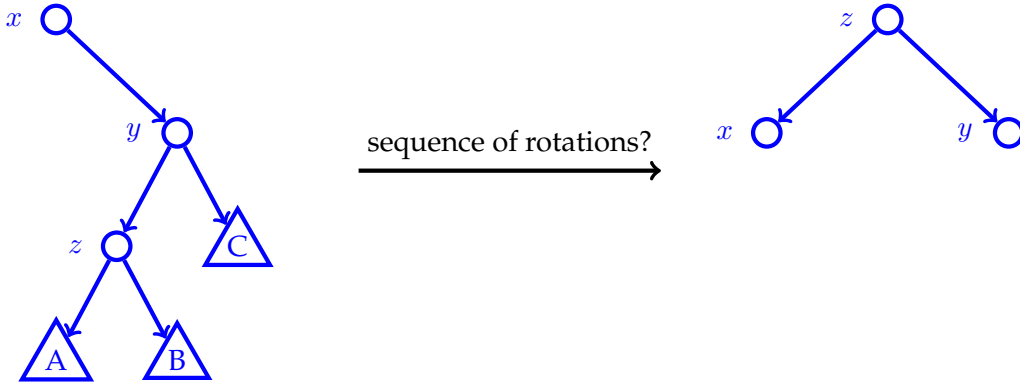
3 Binary Search Trees (30 points)

In this problem we consider plain binary search trees with of `int` keys, without rebalancing.

Task 1 (10 pts). Give a binary search tree with 3 elements (`int` keys) for which there is only one node at which a new `int` key can be inserted. That is, any key inserted into the binary search tree will end up in exactly the same node (assuming that it is not a duplicate of a key already in the tree). Of course, more than one key could be inserted into the binary search tree, but they all need to end up in the same node of the binary search tree.

Task 2 (5 pts). Assume that we have an array sorted *in reverse* from which we want to construct a binary search tree. Which order of insertion should we follow to obtain a tree that is as balanced as possible?

Task 3 (10 pts). The binary search tree on the left has been rotated by a sequence of single left or right rotations to the binary search tree on the right. Unlike the nodes x, y, z , which represent int keys, A, B, and C represent subtrees of unknown size. Complete the shape of the tree on the right by putting the subtrees A, B, C where they belong.



Task 4 (5 pts). Give a sequence of single left or right rotations, and the two nodes being rotated, to transform the binary search tree on the left to the binary search tree on the right that you identified in Task 3. For full credit, find the shortest sequence (which requires less than 5 steps). Indicate in the rightmost column what we know about whether the tree *after* the rotation satisfies the balance invariant of the AVL trees (i.e., the heights of the left and right children of any node differ by at most one) by indicating yes/no/maybe.

Step	Left or Right?	at nodes	AVL?
1			
2			
3			
4			
5			