
Efficient Algorithms and Data Structures I

*Deadline: December 3, 2018, 10:15 am in the **Efficient Algorithms** mailbox.*

Homework 1 (3 Points)

In an ancient book on elephant mathematics, the elephant Amadeus Cage, discovers Algorithm 1. The description has almost faded, but seems to be `elephantnutco`, its subroutine being called `wicked`. The description also indicates that the algorithm takes a boolean array A and an integer k as input.

Algorithm 1: `elephantnutco(A, k)`

```
1 Algorithm elephantnutco(A, k)
2   | for  $j = 0; j < k; j \leftarrow j + 1$  do
3   |   | wicked(A)

1 Procedure wicked(A)
2   |    $i \leftarrow 0$ 
3   |   while  $i < A.length$  and  $A[i] == TRUE$  do
4   |     |  $A[i] \leftarrow FALSE$ 
5   |     |  $i \leftarrow i + 1$ 
6   |   if  $i < A.length$  then
7   |     |  $A[i] \leftarrow TRUE$ 
```

1. Describe briefly what the procedure `wicked` does.
2. Using a suitable potential function, show that the amortized running time of procedure `wicked` during an execution of `elephantnutco` is $\mathcal{O}(1)$, **if the array A is initially false at every position.**

Homework 2 (7 Points)

A minqueue is a dynamic data structure that supports the following operations:

- `ENQ(x)`: Inserts the number x into the minqueue.
- `DEQ()`: Removes the element that has been in the minqueue for the longest time.
- `FMin()`: Returns the smallest value in the minqueue but does NOT remove it.

You may assume that all numbers that are enqueued are distinct.


You are given two doubly linked lists (called A and B) to store elements, as well as some constant amount of additional storage. Describe how to implement a minqueue such that all operations are supported in amortized constant time.

Describe precisely how the operations are implemented. Then analyze their running times using the accounting method or a suitable potential function. You do not need to prove that the operations work correctly.


Hint: Try to implement DEQ() and FMin() in worst-case constant time. Use list A to store all elements in your data structure. Use list B to store all elements that can still become minimum.

Homework 3 (6 Points)

A *Bavarian Search Tree* is a Binary Search Tree in which each node v contains a *base* as well as an additional value called *addend*. The addend of a node v is implicitly added to each base in the subtree rooted at v . Let $(base, addend)$ denote the contents of any node v . For example, the following tree contains the elements $\{5, 6, 7\}$:



(a) In the following Bavarian Search Tree, write the value of each node next to it. The value of the root is 10.



(b) Let h be the height of some Bavarian Search Tree. Describe how to perform the following operations in $O(h)$ time:

- FIND(x): return YES if value x is stored in tree T .
- INSERT(x): inserts element x in tree T .
- SUPERADD(x,k): add k to all elements $\geq x$ currently in the tree.

(c) Show how to perform a right-rotation in a Bavarian Search Tree in constant time.

For parts (b) and (c), show that your implementation reaches the required time bounds. You do not need to prove correctness.

Homework 4 (4 Points)

Suggest how to use a skip list so that given a pointer to a node with key x , we can return a pointer to a node with key $y < x$ in $O(\log k)$ expected time where k is the distance between the nodes with values y and x in L_0 . Prove that your method works!

Tutorial Exercise 1

Prove that there exists a sequence of n insert and delete operations on a $(2, 3)$ -tree s.t. the total number of split and merge operations performed is $\Omega(n \log n)$.

Tutorial Exercise 2

Amanda is developing the app `PrintNeighbors`, that allows users to share their printer with others for a small fee. She decides to use her favorite data structure, Skip Lists, to manage the users. Her boss, however, despises probability theory and asks her to *derandomize* the skip list.

Amanda first tries to design a 1-Skiplist, defined as follows. List L_0 contains all elements. For $k \geq 1$, every second element in L_{k-1} is skipped in L_k , thus dividing the number of elements by 2.

1. Analyze the running time of a search in a 1-Skiplist. Show that inserts cannot be done efficiently in a 1-Skiplist.
2. Amanda wants to improve her data structure to support efficient inserts. Describe a deterministic data structure based on skip lists, that can perform searches and inserts in $\mathcal{O}(\log n)$.

Hint: Think about (a, b) -trees!

Tutorial Exercise 3

Describe a dynamic data structure for storing intervals. Your data structure should support the following operations:

- `INSERT($[q_\ell, q_r]$)`: Stores the interval with left endpoint q_ℓ and right endpoint q_r into the data structure. Returns a handle I for referencing the interval.
- `DELETE(I)`: Deletes the interval referenced by I from the data structure.
- `OVERLAP()`: Returns true if at least two intervals in the data structure overlap and false otherwise.

You may assume that all endpoints of intervals are distinct. Your data structure should perform `INSERT` and `DELETE` in $\mathcal{O}(\log n)$ time and `OVERLAP` in constant time, where n is the number of intervals currently stored in your data structure.

Describe precisely how your data structure is implemented and prove correctness and running time of all operations.

Hint: Use an augmented red-black tree that stores the intervals sorted by their left endpoint.

Hint 2: Observe that `OVERLAP` must take constant time. This limits the possibilities for augmenting the tree.

The more you think about what the B could mean, the more you learn about B-Trees, and that is good.

- R. Bayer