Technische Universität München
Fakultät für Informatik
Lehrstuhl für Effiziente Algorithmen
Prof. Dr. Harald Räcke
Chintan Shah

Wintersemester 2014-15
Assignment 6
November 17, 2014

Efficient Algorithms and Datastructures I

Question 1 (10 Points)

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 $\Omega(n \log n)$.

Question 2 (10 Points)

Show how to maintain a dynamic set Q of numbers that supports the operation MIN-GAP, which gives the magnitude of difference of the two closest numbers in Q. For example, if $Q = \{1, 5, 9, 15, 18, 22\}$, then MIN-GAP($Q$) returns 18-15=3, since 15 and 18 are the two closest numbers in Q. Make the operations INSERT, DELETE, SEARCH, and MIN-GAP as efficient as possible, and analyze their running times.

Question 3 (10 Points)

Suppose that we wish to keep track of a *point of maximum overlap* in a set of itervals - a point that has the largest number of intervals in the set of intervals overlapping it.

1. Show that there will always be a point of maximum overlap which is an endpoint of one of the segments.

2. Design a data structure that efficiently supports the operations INSERT, DELETE, and FIND_POM which are defined as follows:

   (a) INSERT(i, j): Inserts the interval $[i, j]$ in the set of intervals.

   (b) DELETE(i, j): Deletes the interval $[i, j]$ from the set of intervals.

   (c) FIND_POM: Returns a point of maximum overlap.

   (*Hint:* Keep a red-black tree of all the endpoints. Associate a value of $+1$ with each left endpoint, and associate a value of -1 with each right endpoint. Augment each node of the tree with some extra information to maintain the point of maximum overlap.)

Question 4 (10 Points)

Suppose 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.