



CS302 - Assignment 10

Due: Tuesday, March 19 at the beginning of class
Hand-in method: paper



<http://www.stupidinventor.com/2010/03/14/comic-55-happy-pi-day/>

1. [3 points] Given the B-Tree below with $t = 2$, draw the tree that results after inserting 'P', 'A', 'B':



2. [8 points] We saw HEAPSORT which takes the data points, inserts them into a heap and then repeatedly calls EXTRACTMAX to sort the data (recall this was $O(n \log n)$). A similar sorting algorithm exists for binary search trees called TREESORT: call INSERT on each data item to build a binary search tree, then do an inorder traversal of the data to get the sorted list.
 - (a) [3 points] Does it work, i.e. is it guaranteed to always give you data in sorted order? Give a counterexample or a *brief* justification of why it's correct.
 - (b) [3 points] What are the best, average and worst case running times for this algorithm assuming you randomize your data before inserting.
 - (c) [2 points] How do these run-times change if we are using a balanced binary search tree (e.g. a red-black tree)?
3. [5 points] Suppose we decide to speed up our B-TREE-SEARCH algorithm and use binary search within a node rather than a linear search. Show that this make the CPU time required $O(\log n)$ independently of how t is chosen.

4. [8 points] We'd like to add some additional functionality to our B-tree. For each of the functions below, describe succinctly (or write pseudo-code) how to accomplish this and state the worse-case running time.
- (a) [3 points] Find the minimum key in a max B-Tree
 - (b) [5 points] Find the successor of a given key. You can assume that you already have a reference to the key (i.e. you don't have to search for it).

Just for fun

Try implementing your function from Problem 4 and see what values you get for larger n .