Section exercise (12 Feb 2003)

B+Tree

1) Bulk load the B+Tree with values 10, 23, 29, 30, 34, 40, 46, 49, 54, 59, 70, 75
2) Insert 80
3) Insert 24
4) Remove 70
5) Go back to the original bulk loaded B+Tree (step 1). Remove 70
6) Insert 24
7) Insert 80
8) Insert 35
9) Insert 85
10) Delete 59
11) Delete 54
12) How many I/O to find 80?
13) How many I/O to find out 79 is not in index?
14) How many I/O to find 40?
15) How many I/O to find 10?

External Sorting

Suppose that you just finished inserting several records into a heap file and now want to sort those records. Assume that the DBMS uses external sort and makes efficient use of the available buffer space when it sorts a file. Here is some potentially useful information about the newly loaded file and the DBMS software available to operate on it:

The number of records in the file is 4500. Each record is a total of 48 bytes long. The page size is 512 bytes. Each page has 12 bytes of control information on it and the remaining 500 bytes are used to store data. Four buffer pages are available. Each record is not allowed to span more than one page. Assume that quick sort is used.

1) If each record is not allowed to span more than one page, how many records can fit into a page? How many pages are needed to store 4500 records?
2) How many sorted subfiles will there be after the initial pass of the sort, and how long will each subfile be?
3) How many passes (including the initial pass just considered) are required to sort this file?
4) What is the total I/O cost for sorting this file?
5) What is the largest file, in terms of the number of records, you can sort with just four buffer pages in two passes? How would your answer change if you had 257 buffer pages?
6) Extra exercise after section: Repeat (2-5) using tournament sort. Come to office hours for solutions.