

Name: _____

Questions 1–7 are due on the 21st; the rest on the 28th. Please refer to the previous assignment for important instructions on the allowable use of resources and the requirements for electronic and hardcopy submission. Use the newsgroup for questions, clarifications, and general discussion related to this homework and class. Ensure that everything you submit is neat and easy to comprehend; otherwise, it will be ignored, with zero credit. You are encouraged to answer the three ★ questions for extra credit but do note that those questions are graded much more strictly than the others.

1. (1 pt.) Write your name in the space provided above.
2. (1 pt.) The pen-and-paper questions in this assignment should be answered on separate sheets of paper that are neatly attached to your hardcopy submission with a single staple in the top left corner. Note the number of additional pages attached: _____
3. (18 pts.) A k -tree, for integer $k > 1$, is either *empty*, denoted \perp , or has a *root* that is either (a) an *l-node*, denoted \circ , or (b) an *i-node*, denoted $r = (t_1, t_2, \dots, t_c)$, comprising a sequence of c subtrees t_i , for some $1 \leq c \leq k$, which are all nonempty k -trees. In the last case, the roots of these c nonempty subtrees are the *children* of r , and r is their *parent*. Depict all distinct (i.e., nonisomorphic) 4-trees with n nodes, for $n = 0, 1, 2, 3, 4, 5$. (A *node* is either an l-node or an i-node.)
4. (10 pts.) A *labeled k -tree* is a k -tree in which each l-node n has a *label* $l(n) \in \mathcal{D}$, for some *label-domain* \mathcal{D} .

Indicate the number of distinct labeled 4-trees with n nodes, for $n = 0, 1, 2, 3, 4, 5$, with $\mathcal{D} = \{1, 2, 3\}$. A tree need not contain all the labels from the domain, and labels may be repeated. You do *not* need to depict all the trees, but you must explain clearly how you arrive at the answers. (Depict at least a few representative trees.) There is no credit for answers without proper explanations.

5. (20 pts.) A *k -search-tree* is a labeled k -tree with a *totally ordered* label-domain, and the following constraint on each i-node n :

Let x_1, x_2, \dots, x_c be the subtrees of n , and let $L(x_i)$ denote the set of labels in subtree x_i . Then, for $i = 1, 2, \dots, c$, every label in $L(x_i)$ is greater than every label in $L(x_{i-1})$ and less than every label in $L(x_{i+1})$, where we define $L(x_0) = \{-\infty\}$ and $L(x_{c+1}) = \{\infty\}$ for convenience.

Depict all distinct labeled 4-search-trees with n nodes, for $n = 0, 1, 2, 3, 4, 5$, with labels drawn from the domain $\{1, 2, 3, 4, 5\}$. You may abbreviate depictions *provided* the resulting depictions are unambiguous. Indicate the total number of trees.

6. (10 pts.) Depict a k -search-tree in which some label l occurs in two distinct locations, or explain clearly why no such tree exists.
7. (30 pts.) Describe efficient algorithms for the following, using the k -search-trees defined earlier. For each algorithm, provide (1) a clear English description, (2) detailed pseudocode, and (3) a proof that the algorithm's running time is $O(h)$ where h is the height of the k -search-tree.
 - (a) Locating a given label in the tree or determining that it is absent from the tree.
 - (b) Inserting a given label into the tree or determining that it is already in the tree.
 - (c) Removing a given label from the tree or determining that the tree does not contain that label.
8. (10 pts.) Conduct an experimental evaluation of the three algorithms for the maximum contiguous subsequence sum problem described in the textbook.¹ Ensure that your results are meaningful, given the limited clock accuracy and other factors. In your electronic submission, include your source code as usual, along with a *plain text ASCII* file that tabulates the running times for input sequences of varying lengths. Describe in that file how the input sequences are generated, and why that generation process is suitable. (Example: Do not exclusively use sequences that contain only positive or only negative numbers.) Explain any discrepancies between your results and those predicted by the textbook. You may use others' code with proper attribution. Also include in your submission a *portable PDF* [sic] file that depicts the running time data in a suitable graphical form.
9. (90 pts.) Implement your algorithms of Question 7 as part of the simple record-manager application described below. Package and submit your source code as in the previous assignment, being sure to include adequate documentation. Poorly written or poorly documented code is likely to receive zero credit.

The record manager stores *key-data pairs* of the form (k, d) where k is a non-negative integer and d is a string. Your application must support *arbitrarily large* integers and strings. The set of key-data pairs must be stored in a labeled k -tree, with each pair stored as the label of a suitable l -node. All user interaction with the record manager is through a text-mode command language based on the standard-input/standard-output interface. The input consists of one command per line. The record manager reads and responds to each command in turn. Except for the `xp` command, the response to each command is also a single newline-terminated line. The syntax and semantics of the commands, and the desired outputs, are outlined by Table 1.

Ensure that (1) your program produces exactly the output described above, with no spurious text such as extra spaces or newlines, prompts, or other messages and (2) your program reads from standard input and writes to standard output, with no additional assumptions on either.

¹Mark Allen Weiss, *Data Structures and Problem Solving Using Java*, 4th edition (Addison-Wesley, 2010), §5.3.

command	actions
<code>c_ k</code>	Create a record manager that contains no records and that uses a k -search-tree with the given value of k . The output is an empty line (single newline character).
<code>s_ k_ d</code>	The pair (k, d) is stored in the record manager. If a pair of the form (k, d') already exists, then the new pair replaces it. The output is an empty line.
<code>e_ k</code>	If the record manager contains the key k then the output is 1 else the output is 0.
<code>r_ k</code>	The data d associated with key k is retrieved from the record manager. The output is d . If the key k is absent from the record store, the output consists of an empty line.
<code>d_ k</code>	If there is a record with key k then it is deleted else there is no effect on the tree. The output is an empty line.
<code>xs</code>	The output is the size of the tree (a single integer).
<code>xh</code>	The output is the height of the tree.
<code>xp</code>	The output consists of the following <i>two-dimensional text representation</i> of the current state of the tree used to store the records. This representation of a tree t contains one line for each node of t . The line representing a node n at depth d begins with $3d$ spaces. If the node is an l-node then these spaces are followed by n 's (k, d) pair, printed as $(k_ d)$; otherwise, they are followed by the integer c , the number of n 's children. Further, the line representing a node is printed before those of its descendants, and all the lines representing the descendants of a node precede the line representing that node's right sibling, if any.

Table 1: Command syntax for the record-manager application. The symbol `_` denotes a single space character.

10. (10 pts.) Conduct an experimental evaluation of each command of your record-manager implementation, following the guidelines of Question 8.
11. (20 pts.) ★ Describe an efficient algorithm for generating all k -search-trees with n nodes, with labels drawn from L , where k , n , and L are provided as input. Provide an English description and pseudocode. Explain why the algorithm is correct. Determine and prove its running time. Implement the algorithm and conduct an experimental study of the running time. Package and submit your implementation as in earlier questions.
12. (20 pts.) ★ Describe modifications to the algorithms of Question 7 that ensure that the height of a k -search-tree with n keys is $O(\log n)$. Follow the instructions of Question 11.
13. (10 pts.) ★ Provide an efficient algorithm for pretty-printing a k -search-tree based on conventional graphical notation using ASCII graphics. Use the newsgroup for elaboration of this question and document and package your results as in Question 11.

Guidelines and reminders

- This homework requires substantial effort. There are many challenges that (by design) will not become apparent until some previous work has been performed. So:

Please start work on both parts right away.

If you start early, you are likely to enjoy the challenges. If you delay then the challenges will turn into frustrations.

- Refer to the class newsgroup for questions, hints, and additional guidelines.
- Coding convention:
 - Do not define packages for your homework code.
 - Format your code to fit well in 80-character lines.
 - Use descriptive names for variables, classes, etc.
 - Follow Java conventions for cases of names (e.g., MyClass, aMember).
- Test your code well, outside of any IDE you may use.
- Submit source code and only source code.
- If you use others' code, be sure to provide proper attributions and to very clearly delineate your code (or modifications) from that code. You should not use others' code if it answers the main part of any question. (You may use it for supplementary tasks, but not the main task.)