CSE 2320   Name _____

Test 1

Fall 2003   Student ID # _____

Multiple Choice. Write your answer to the LEFT of each problem. 4 points each

1. PARTITION is useful in solving all of the following problems, except
    A. Finding the k smallest elements
    B. QUICKSORT
    C. Selection
    D. Splitting an array for MERGESORT
2. Which of the following sorts is not stable?
    A. INSERTION-SORT
    B. LSD RADIX-SORT
    C. MERGESORT
    D. QUICKSORT
3. Which of the following sorts is not based on key comparisons?
    A. INSERTION-SORT
    B. LSD RADIX-SORT
    C. MERGESORT
    D. QUICKSORT
4. Suppose the input to HEAPSORT is always a table of identical integers. The worst-case time will be
    A. $\Theta(1)$
    B. $\Theta(n)$
    C. $\Theta(n \lg n)$
    D. $\Theta(n^2)$
5. Suppose the input to PARTITION is always a set of identical integers. The worst-case time will be
    A. $\Theta(1)$
    B. $\Theta(n)$
    C. $\Theta(n \lg n)$
    D. $\Theta(n^2)$
6. The time to multiply two n x n matrices is:
    A. $\Theta(n)$
    B. $\Theta(n \lg n)$
    C. $\Theta(n^2)$
    D. $\Theta(n^3)$
7. $f(n) = n \lg n$ is in all of the following sets, except
    A. $\Omega(\lg n)$
    B. $\Theta(\log(n!))$
    C. $O(1/n)$
    D. $O(n^2)$
8. Which of the following best approximates $H_m - H_n$? (m > n)
    A. H_{m-n}
    B. $1/(m - n)$
    C. $\ln(m / n)$
    D. $\ln(m - n)$
9. Which sort may be viewed as being an improvement on insertion sort?
    A. COUNTING-SORT
    B. HEAPSORT
    C. QUICKSORT
    D. shellsort
10. Counting sort is often used as a part of which sorting algorithm?
    A. INSERTION-SORT
    B. LSD RADIX-SORT
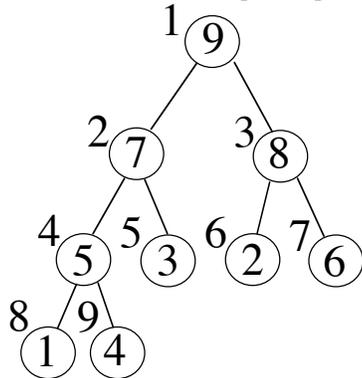   C. MERGESORT
    D. QUICKSORT

Long Answer

1. Prove that if $f(n) \in O(g(n))$ then $\dfrac{1}{f(n)} \in \Omega\left(\dfrac{1}{g(n)}\right)$. 10 points

2. Use the substitution method to show that T(n) = 2T(n/4) + n is in $\Theta(n)$. 15 points

3. Use the recursion-tree method to show that $T(n) = 2T(n/4) + n$ is in $\Theta(n)$. 15 points
4. Demonstrate PARTITION on the following array. 10 points

$$9 \quad 3 \quad 6 \quad 2 \quad 7 \quad 4 \quad 1 \quad 8 \quad 0 \quad 5$$

5. Show the maxheap after performing HEAP-EXTRACT-MAX two times. 10 points

CSE 2320   Name _____

Test 2

Fall 2003   Student ID # _____

Multiple Choice. Write the letter of your answer to the LEFT of each problem. 4 points each

Problems 1 and 2 refer to the following hash table whose keys are stored by linear probing using $h(key) = key \% 13$.

0	1	2	3	4	5	6	7	8	9	10	11	12
			94		122	110	20	86	87	62		

1. 143 would be inserted into which slot of the given table?
   A. 0
   B. 1
   C. 2
   D. 11
2. 136 would be inserted into which slot of the given table?
   A. 0
   B. 4
   C. 6
   D. 11

Problems 3 and 4 refer to the following hash table whose keys are stored by double hashing using $h_1(key) = key \% 13$ and $h_2(key) = 1 + (key \% 12)$.

0	1	2	3	4	5	6	7	8	9	10	11	12
			120	186	187	162		122	110			194

3. 263 would be inserted into which slot of the given table?
   A. 0
   B. 2
   C. 3
   D. 7
4. 303 would be inserted into which slot of the given table?
   A. 0
   B. 2
   C. 3
   D. 7
5. Why is it common for a circular queue implementation to waste one table element?
   A. To avoid confusing an empty table with a full table
   B. To have a place to store the tail and head values
   C. To make sure the queue always has at least one element in use
   D. To perform some loops faster
6. When evaluating a postfix expression, the stack contains
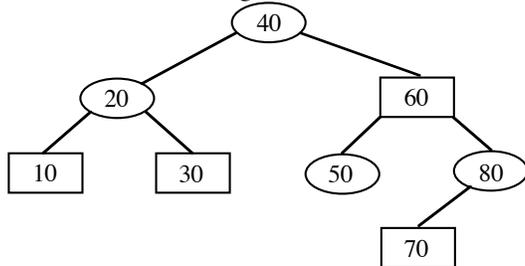   A. Both operands and operators
   B. Both parentheses and operators

C. Operands only
D. Operators only
7. The worst-case time to find the maximum key in a circular, doubly-linked list with n nodes in ascending order is:
   A. $\Theta(1)$
   B. $\Theta(\log n)$
   C. $\Theta(n \log n)$
   D. $\Theta(n)$
8. How should the successor of a leaf in an unbalanced binary search tree be found?
   A. Examine the ancestors of the leaf
   B. Go right, then proceed to the left
   C. Inorder traversal
   D. Preorder traversal
9. If POP is implemented as return stack[--SP], then PUSH of element X is implemented as:
   A. return stack[SP++]
   B. stack[SP++] = X
   C. stack[--SP] = X
   D. stack[++SP] = X
10. Which of the following is not true about probe sequences for an implementation of double hashing?
    A. All slots in the hash table appear in each probe sequence
    B. Every key has a probe sequence different from the probe sequences for other keys
    C. The elements of a probe sequence are subscripts in the hash table
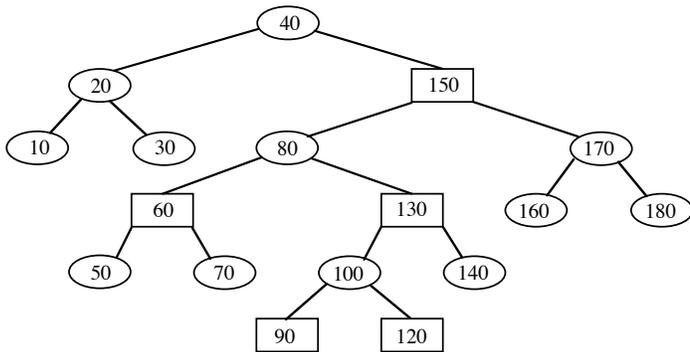   D. The probe sequence for a key cannot change

Long Answer
1. Identify the problems with the following attempt at constructing a red-black tree. (10 points)
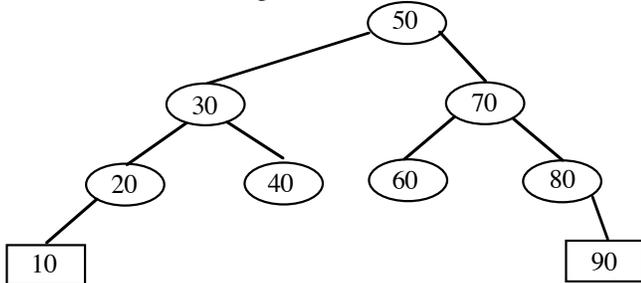

2. Determine (analytically) the expected unsuccessful search performance of the following data structures when 5,000,000 records are stored. You may assume that the keys are equally likely to be requested. (10 points)
   a. Chaining with a table with 500,000 entries. Each linked list is unordered.
   b. Double hashing with a table with 7,500,013 entries. (7,500,013 is prime.)
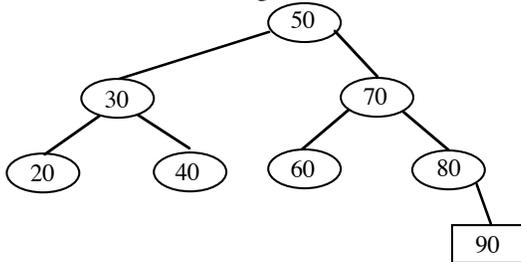3. Insert 75 into the given red-black tree. Be sure to indicate the cases that you used. (10 points)

4. Insert 110 into the given red-black tree. Be sure to indicate the cases that you used. (10 points)

40
20   150
10 30   80   170
60   130   160 180
50 70 100 140
90 120

5. Delete 40 from the given red-black tree. Be sure to indicate the cases that you used. (10 points)

50
30   70
20   40   60   80
10   90

6. Delete 30 from the given red-black tree. Be sure to indicate the cases that you used. (10 points)

50
30   70
20   40   60   80
90

CSE 2320   Name _____
Test 3
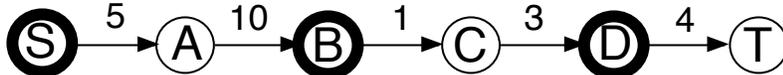Fall 2003   Student ID # _____
Multiple Choice. Write the letter of your answer to the LEFT of each problem. 2 points each
1. Before searching for a minimum cut in a network, it is useful to do the following:
   A. Determine the type of each edge using depth-first search.
   B. Find and record augmenting paths until none remains.
   C. Find one augmenting path.
   D. Perform a breadth-first search on the input network.
2. A fail link of -1 requires the KMP matcher to take what action?
   A. Give up the search entirely, since the pattern cannot appear within the text.
   B. Move both pointers up one symbol.
   C. Move the pattern pointer to the next pattern symbol and set the text pointer to 0.
   D. Move the text pointer to the next text symbol and set the pattern pointer to 0.
3. Compressed adjacency lists have the following disadvantage:
   A. Testing whether an edge from X to Y is present will take $\theta(V + E)$ worst-case time.
   B. They are static.
   C. They can only be used for graphs without weights.
   D. They require $\theta(V + E)$ space to store.
4. During a breadth-first search, the status of a gray vertex is:
   A. It has been completely processed.
   B. It is in the FIFO queue.
   C. It is in the priority queue.
   D. It is undiscovered.
5. The capacity of any cut is:
   A. A lower bound on the maximum flow.
   B. An upper bound on the maximum flow.
   C. The same as the capacity of all other cuts.

D. The same as the maximum attainable flow.
6. Augmenting paths are usually found by:
   A. Breadth-first search on the saturated edges in the residual network.
   B. Breadth-first search on the unsaturated edges in the residual network.
   C. Depth-first search on the unsaturated edges in the residual network.
   D. Dijkstra's algorithm.
7. Which of the following is true about KMP string search?
   A. Once the fail links have been constructed, the pattern is no longer needed.
   B. The fail links are constructed based on the pattern and may be applied to different texts.
   C. The fail links are constructed based on the text and may be applied to different patterns.
   D. The fail links are constructed for a particular pattern and a particular text.
8. An adjacency matrix is the most useful representation for which problem?
   A. Breadth-first search
   B. Finding strongly-connected components
   C. Maximum network flow
   D. Warshall's algorithm
9. The number of HEAP-EXTRACT-MINs to build a Huffman code tree for n symbols is:
   A. $\theta(\log n)$
   B. n - 1
   C. n
   D. 2n - 2
10. During depth-first search on a directed graph, a cycle is indicated by which edge type?
    A. Back
    B. Cross
    C. Forward
    D. Tree
11. Which of the following is a longest common subsequence for 0 1 2 0 1 2 and 0 0 1 1 2 2?
    A. 0 0 1 1
    B. 0 0 1 1 2
    C. 0 0 1 2
    D. 0 1 2 0
12. The worst-case time for depth-first search is:
    A. $\theta(V + E)$
    B. $\theta(E \lg V)$
    C. $\theta(V \lg V)$
    D. $\theta(V \lg E)$
13. Prim's algorithm, when implemented with a heap, is most suitable for:
    A. Finding the minimum spanning tree of a dense graph.
    B. Finding the minimum spanning tree of a sparse graph.
    C. Finding the shortest paths from a designated source vertex in a dense graph.
    D. Finding the shortest paths from a designated source vertex in a sparse graph.
14. The Edmonds-Karp variant is important because:
    A. It solves the bipartite matching problem.
    B. It solves the network flow problem in polynomial time.
    C. It solves the network flow problem using critical edges.
    D. It solves the network flow problem without using augmenting paths.
15. The capacity of the following cut is _____. (S vertices are bold.)

    A. 1
    B. 10
    C. 13
    D. 23
16. The fastest method for finding the diameter of a tree is to:
    A. Use breadth-first search.
    B. Use Dijkstra's algorithm.
    C. Use the Floyd-Warshall algorithm.
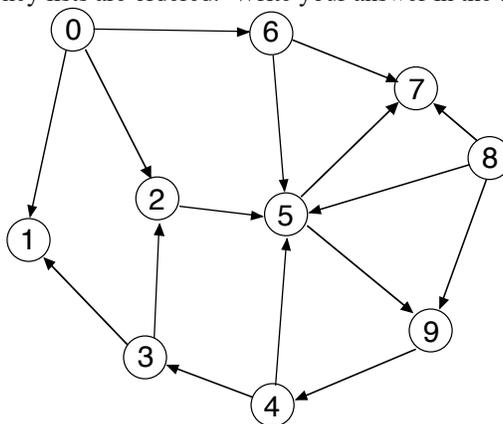   D. Use the Ford-Fulkerson algorithm.

17. Suppose that a depth-first search on a directed graph yields a path of tree edges from vertex X to vertex Y. If there is also an edge from X to Y, then its type will be:
    A. Back
    B. Cross
    C. Forward
    D. Tree
18. Suppose that there is only one path from vertex 5 to vertex 10 in a directed graph:
    5 → 7 → 8 → 3 → 2 → 10. During the scan of which column will Warshall's algorithm record the presence of this path?
    A. 2
    B. 5
    C. 8
    D. 10
19. When finding the strongly connected components, the number of components is indicated by:
    A. The number of back edges found during the first depth-first search.
    B. The number of cross edges found during the second depth-first search.
    C. The number of restarts for the first depth-first search.
    D. The number of restarts for the second depth-first search.
20. Which of the following is solved heuristically by a greedy method?
    A. Fractional knapsack
    B. Finding the shortest paths from a designated source vertex in a sparse graph.
    C. Minimum spanning tree
    D. 0/1 knapsack

Long Answer
1. Fill in the KMP failure links. 10 points.

	1	pattern	2
0		a	
1		b	
2		c	
3		d	
4		a	
5		b	
6		c	
7		a	
8		b	
9		a	

2. Perform depth-first search on the following graph, including start/finish times and edge types (T=tree, B=back, C=cross, F=forward. Assume that the adjacency lists are ordered. Write your answer in the tables below. 15 points
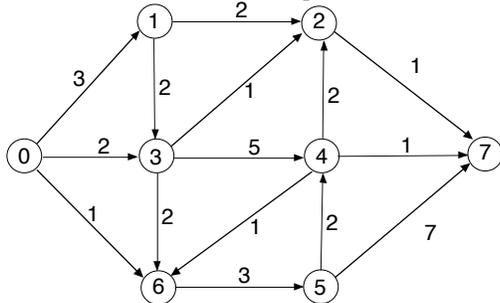

Vertex	Start	Finish	Edge	Type	Edge	Type
0	1	___	0 1	___	6 5	___
1	___	___	0 2	___	6 7	___
2	___	___	0 6	___	8 5	___
3	___	___	2 5	___	8 7	___
4	___	___	3 1	___	8 9	___
5	___	___	3 2	___	9 4	___
6	___	___	4 3	___		
7	___	___	4 5	___		

3. 15 points - 5 points for each part.

a. Suppose that A is a binary adjacency matrix for a directed graph with n vertices (numbered 0 to n - 1). Give code for Warshall's algorithm. Input, output, and recording successors are NOT needed.

b. Suppose that C is an adjacency matrix that stores the capacities for a network with n vertices. Give code to modify C to record the capacity of a maximum capacity path between each pair of vertices. Input, output, and recording successors are NOT needed.

c. The following successor matrix was produced by Warshall's algorithm. Give the path from vertex 0 to vertex 4.

```
-1    3    3    3    3
-1    3    3    3    4
-1    1    1    1    4
-1    2    2    2    2
-1   -1   -1   -1   -1
```

4. Give augmenting paths for determining a maximum flow and give a minimum cut for the following network. 0 is the source and 7 is the sink. 10 points.

S vertices: 0
T vertices: 7
Augmenting Paths and Contribution to Flow:

5. Complete the following instance of the optimal matrix multiplication ordering problem, including the tree showing the optimal ordering. 10 points

```
p[0]=5
p[1]=2
p[2]=2
p[3]=4
p[4]=6
```

	1		2		3		4	
1	0	0	20	1	56	1	???	?
2	-------		0	0	16	2	64	3
3	-------		-------		0	0	48	3
4	-------		-------		-------		0	0