



EXAMINATIONS — 2010

END-OF-YEAR

COMP 261
ALGORITHMS
and
DATA STRUCTURES

Time Allowed: 3 Hours

Instructions: Attempt ALL Questions.

Answer in the appropriate boxes if possible — if you write your answer elsewhere, make it clear where your answer can be found.

The exam will be marked out of 180 marks.

Non-programmable calculators without a full alphabetic key pad are permitted.

Non-electronic foreign language dictionaries are permitted.

Useful formulas are listed on the last page of the exam.

Questions	Marks
1. File Structures	[22]
2. Union Find	[17]
3. Searching in Graphs	[26]
4. String Searching	[16]
5. Graphics Rendering	[9]
6. Advanced Tree Structures	[30]
7. Text Processing	[40]
8. File Structures and Hashing	[20]

Question 1. File Structures

[22 marks]

Suppose a file contains 100,000 fixed length records describing individual students. Each record has the following fields:

StudentID: (length = 5 characters),

Name: (length = 50 characters),

Address: (length = 120 characters),

DoB: (length = 10 characters).

Assume that the file blocks are stored contiguously and that the block size for the file is 1024 characters.

(a) [4 marks] Calculate the record size L in characters. Show your working.

(b) [4 marks] Calculate the blocking factor f and the number of file blocks b . Assume an unspanned file organisation. Show your working.

(c) [4 marks] Calculate the *average* number of block accesses needed to perform a linear search for a random record in the file given its StudentID. Show your working.

(Question 1 continued on next page)

(Question 1 continued)

(d) [4 marks] Calculate the *worst case* number of block accesses needed to perform a binary search for a random record in the file given its StudentID. Assume the file is ordered by StudentID. Show your working.

(e) [6 marks] Explain the differences between primary file organisation and secondary file organisation.

Question 2. Union Find

[17 marks]

(a) [4 marks] Explain why Kruskal's algorithm for finding minimum spanning trees needs to keep track of a set of sets of nodes.

An Efficient Union-Find algorithm:

MakeSet(x):

```
x.parent ← x
x.rank ← 0
add x to collection of sets.
```

Find(x):

```
if x.parent = x then return x
else
  x.parent ← Find(x.parent)
return x.parent
```


Union(x, y):

```
xroot ← Find(x)
yroot ← Find(y)
if xroot.rank < yroot.rank then
  xroot.parent ← yroot
  remove xroot from collection of sets.
else
  yroot.parent ← xroot
  remove yroot from collection of sets.
if xroot.rank = yroot.rank then
  xroot.rank++
```

(Question 2 continued on next page)

(Question 2 continued)

(b) [8 marks] Consider the following diagram of a collection of four sets represented using the Union-Find data structure. The numbers are the ranks of the nodes.



Using the efficient Union-Find algorithm on the facing page, show the changes that will be made to the graph by the following *sequence* of operations.

- (i) Find(Q)
- (ii) Union(A, B)
- (iii) Union(B, F)
- (iv) Find(J)
- (vi) Union(N,J)

Draw the changes on the diagram above.
 Hint: note the order of the arguments of Union carefully.

(Question 2 continued)

SPARE PAGE FOR EXTRA ANSWERS

Cross out rough working that you do not want marked.
Specify the question number for work that you do want marked.

(Question 2 continued on next page)

(Question 2 continued)

(c) [5 marks] The efficient Union-Find algorithm moves nodes around in the tree when searching for a value in order to make future searches in the tree more efficient. Could a similar technique be used for improving binary search trees? If you think it could, briefly suggest how and the conditions under which it might help. If you think it couldn't, explain why not.

Question 3. Searching in Graphs

[26 marks]

(a) [5 marks] Prim's algorithm for minimum spanning trees is very similar to Dijkstra's algorithm for single source shortest paths. Explain the key way in which Prim's algorithm differs from Dijkstra's algorithm, and explain why this difference results in two different trees.

(b) [5 marks] Give an example of a small graph, including a start node and a goal node, and an *inadmissible* heuristic estimate of remaining path length for each of the nodes and show how A* would find the wrong path to the goal node using this heuristic. Show the path that A* would find and show the shortest path that it should have found.

Note: show all the edge lengths and show the heuristic estimate on each node.


S/

G/0

(Question 3 continued on next page)

(Question 3 continued)

(c) [8 marks] Suppose you are using A* to search for the shortest path from S to G in the graph below, where the heuristic estimate for each node is shown in the node (the heuristic is admissible). Show the order in which nodes will be *added* to the queue, and the order in which they are *removed* from the queue. (When visiting a node, consider the neighbours of the node in alphabetic order.) Hint: keep track of the queue, along with the total path length (path so far plus heuristic estimate) for all the nodes on the queue.



Added to Queue:


Removed from Queue:

(Question 3 continued on next page)

(Question 3 continued)

(d) [8 marks] Suppose you are using Dijkstra's algorithm to find the shortest path from **S** to **G** in the graph below (the same as in part (c), but without heuristic estimates). Show the order in which nodes will be *added* to the queue, and the order in which they are *removed* from the queue. (When visiting a node, consider the neighbours of the node in alphabetic order.)

Hint: keep track of the queue, along with the priority for all the nodes on the queue.



Added to Queue:

Removed from Queue:

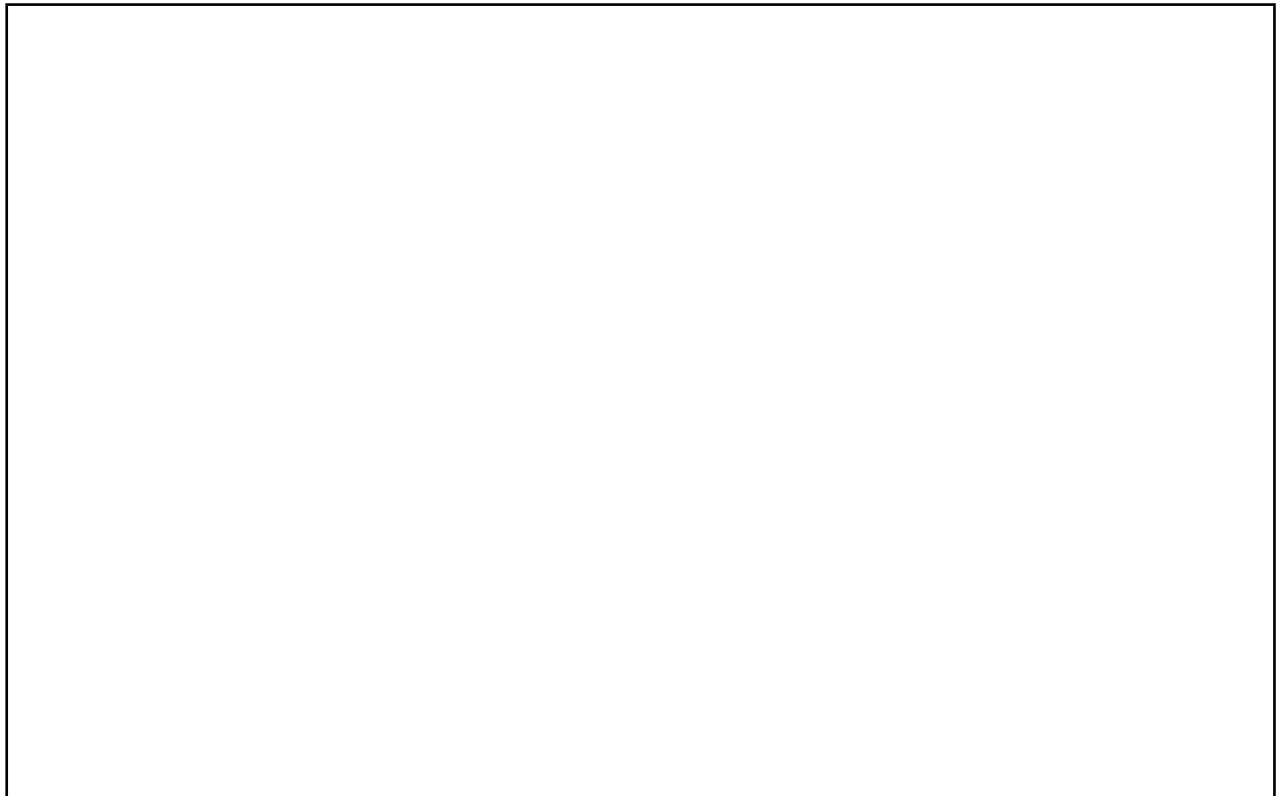
Question 4. String Searching

[16 marks]

(a) [5 marks] To search a text for an occurrence of any of a set of words, it is efficient to construct a trie of the words. Draw a trie for the following set of words.

Note: The characters should be attached to the edges in the trie, and all terminal nodes should be indicated clearly.

red	strike	sell	search	read	sea
ready	strings	strict	reach	seat	string



(Question 4 continued on next page)

(Question 4 continued)

The simplified Boyer-Moore string searching algorithm shown on the facing page constructs a “Bad-Character” table and then uses it while matching to work out how far to move the string forward when there is a mismatch. It does not use a “GoodSuffix” table.

If the string being searched for is agkga, the BadCharacter table would be the following:

BC:

0	5	5	5	5	5	1	5	5	5	2	5	5	5	...
'a'	'b'	c	d	e	f	g	h	i	j	k	l	m	n	...

(b) [11 marks] Show the first 8 sets of values of *i*, *k*, and *T*[*i*] at the point marked *** (the test in the while loop) in the simplified Boyer-Moore algorithm when it is used to search for the string agkga in the following text. (The first set of values is given.)

a g k b b a g k g k k g a g k g a m c d k g a a k a
0 5 10 15 20 25

<i>i</i>	<i>k</i>	<i>T</i> [<i>i</i>]
4	4	b

(Question 4 continued on next page)

(Question 4 continued)

A simplified version of the Boyer-Moore algorithm which only uses the “BadCharacter” table, not the “goodSuffix” table.

Boyer–Moore Search:

Input: *string* S[0 .. m-1],
 text T[0 .. n-1]

Output: the position in T at which S is found, or -1 if not present

Variables: i // position of current character in T
 k // position of current character in S
 BC[*char*] // Bad Character table

Actions:

BC ← computeBadCharTable(S)

i ← m-1

while i < n **do**

 k ← m-1

while k ≥ 0 and S[k] = T[i] **do** *** // record i, k, and T[i] at this point

 i ← i-1,

 k ← k-1

if k = -1 **then return** i+1

 i ← i + max(BC[T[i]], m-k) // mismatch ⇐ advance

return -1

SPARE PAGE FOR EXTRA ANSWERS


Cross out rough working that you do not want marked.
Specify the question number for work that you do want marked.

Question 5. Graphics Rendering

[9 marks]

(a) [6 marks] Show the values in the edge-lists that would be constructed when rendering the following polygon. The (x, y, z) coordinates of the vertices are shown.

(Note that the z values are chosen carefully to make the interpolation easy.)



Edge-Lists

	x_{left}	z_{left}	x_{right}	z_{right}
20				
21				
22				
23				
24				


(b) [3 marks] In constructing and using the edge-lists, you had to convert floating point numbers to integers. Explain why this can introduce errors unless you are careful.

Question 6. B-Trees

[30 marks]

(a) [10 marks] State the constraints that must be maintained by a *B*-tree with order $p = 2m + 1$ and height h . Indicate which constraints ensure that a *B*-tree will be kept well-balanced.

(b) [10 marks] Consider the *B*-tree of order 5 illustrated below.



Update the *B*-tree by successively deleting the key values 75, 15, 62, 21. In your answer, show the *B*-tree after each deletion and briefly describe what you have done.

Note, the empty trees below are to save you time; you may modify their structure if you choose.

The *B*-tree after deleting key value 75:


```

    graph TD
      Root[" | | | "]
      L1[" | | | "]
      L2[" | | | "]
      L3[" | | | "]
      L4[" | | | "]
      Root --> L1
      Root --> L2
      Root --> L3
      Root --> L4
  
```


(Question 6 continued on next page)

(Question 6 continued)


The *B*-tree after deleting key values 75 and 15:



The *B*-tree after deleting key values 75, 15, and 62:



The *B*-tree after deleting key values 75, 15, 62, and 21:



SPARE PAGE FOR EXTRA ANSWERS

Cross out rough working that you do not want marked.
Specify the question number for work that you do want marked.

(Question 6 continued)

(c) [5 marks] Consider a B -tree file of order $p = 2m + 1 = 199$ that contains $r = 10^7$ records. How many disk accesses will it take to retrieve a random record from the file using the B -tree in the worst case?

You may assume that the block size is larger than the size of the B -tree nodes and larger than the size of the records.

Show your working.

(d) [5 marks] Explain the difference between B -tree and B^+ -tree index structures.

Question 7. Text Processing

[40 marks]

(a) [5 marks] Explain the difference between an *Abstract Syntax Tree* and a *Concrete Parse Tree*.



(Question 7 continued on next page)

(Question 7 continued)

(b) [10 marks] Write a parser for the simple Query grammar below that returns `true` or `false`. Assume nonterminals are always in quotes and there are spaces separating tokens for simplicity.

```
QUERY ::= "SELECT" "*" ["FROM" NAME] ["WHERE" NAME "=" DATA] ";"  
NAME  ::= [A-Za-z]+  
DATA  ::= [A-Za-z0-9]*
```

Hint: It is easier to write the parser in Java, but you may use clear pseudocode if you prefer.

```
public boolean parseQuery(Scanner s)
```

(Question 7 continued on next page)

(Question 7 continued)

Consider the following grammar where nonterminals are in uppercase and terminals are enclosed in quotation marks. Assume that tokens will be separated by spaces.

```
EXPRESSION ::= FOO "+" BAR "end"  
FOO ::= [a-z0-9]+ | EXPRESSION  
BAR ::= EXPRESSION FOO | a*b*c+
```

(c) [10 marks] For each of the following sentences, state whether it belongs to the language defined by this grammar.

- | | |
|--------------------------|-----------------------------------|
| <input type="checkbox"/> | abracadabraend + abc end |
| <input type="checkbox"/> | 0 + c end + c end |
| <input type="checkbox"/> | aabbc + aacc end |
| <input type="checkbox"/> | 0 + abc end + c end + c end |
| <input type="checkbox"/> | 0 + 0 + c end c end c end + c end |

(d) [5 marks] Explain why the grammar above cannot be parsed by a predictive, one symbol lookahead, left-to-right (LL(1)) parser.

(Question 7 continued on next page)

(Question 7 continued)

(e) [10 marks] Write a new grammar for EXPRESSIONs which specifies the same language as the grammar on the facing page, but which can be parsed using an LL(1) parser.

EXPRESSION ::=

Question 8. File structures and Hashing

[20 marks]

(a) [10 marks] For each of the following file structures, discuss their advantages and disadvantages by explaining the efficiency of the different file operations (insertion, deletion, search, sequential access) with different structures, and giving examples of when it is appropriate to use each kind of file.

(i) Heap file:

(ii) Sequential file.

(iii) Hash file.

(Question 8 continued)

(b) [10 marks] Describe how extendible hash files work.

SPARE PAGE FOR EXTRA ANSWERS

Cross out rough working that you do not want marked.
Specify the question number for work that you do want marked.

Useful Formulas

You may tear off this page if you wish. You do not need to hand it in.

File Performance Formulas

- blocking factor: $f = \lfloor \frac{B}{L} \rfloor$
- number of blocks: $b = \lceil \frac{r}{f} \rceil$
- external sort-merge: $N = 2b \cdot (1 + \lceil (\log_{n-1} b) - 1 \rceil) = 2b \cdot (1 + \lceil (\frac{\log_{10} b}{\log_{10}(n-1)} - 1) \rceil)$
(where n is the number of buffers)

B-tree (worst case)

- height: $h = 1 + \lfloor \log_{m+1} \frac{r+1}{2} \rfloor = 1 + \lfloor \log_{10} \frac{\frac{r+1}{2}}{\log_{10}(m+1)} \rfloor$
- number of leaves: $N_{leaves} = 2(m+1)^{h-2} \leq N_{leaves} \leq (2m+1)^{h-1}$

B⁺-tree (worst case)

- height: $h = 2 + \lfloor \log_{m+1} \frac{r}{2m} \rfloor = 2 + \lfloor \frac{\log_2 \frac{r}{2m}}{\log_2(m+1)} \rfloor$
- number of leaves: $N_{leaves} = \lceil \frac{r}{m} \rceil$

Index-Sequential File with a B-tree

- number of sequence sets: $s = \lceil \frac{r}{f} \rceil \leq s \leq \lceil \frac{2r}{f} \rceil$

Logs to base 2

n	1	2	4	8	16	32	64	128	256	512	1,024	1,048,576
$\log_2 n$	0	1	2	3	4	5	6	7	8	9	10	20

Logs to base 10

n	5	10	50	100	500	1000	5000	10,000	10 ⁶	5 × 10 ⁶	10 × 10 ⁶	50 × 10 ⁶
$\log_{10} n$	0.7	1	1.7	2	2.7	3	3.7	4	6	6.7	7	7.7