

Computer Science Foundation Exam

May 17, 2025

Section A

BASIC DATA STRUCTURES

**NO books, notes, or calculators may be used,
and you must work entirely on your own.**

SOLUTION

Question #	Max Pts	Category	Score
1	5	ALG	
2	10	DSN	
3	10	ALG	
TOTAL	25	----	

You must do all 3 problems in this section of the exam.

Problems will be graded based on the completeness of the solution steps and not graded based on the answer alone. Credit cannot be given unless all work is shown and is readable. Be complete, yet concise, and above all be neat. For each coding question, assume that all of the necessary includes (stdlib, stdio, math, string) for that particular question have been made.

1) (5 pts) ALG (Dynamic Memory Management in C)

Consider the following code that attempts to allocate a new array with the same number of rows as arr, but triple the columns, copy the values from arr into the new array (in the same slots), free the dynamically allocated memory pointed to by arr and return a pointer to the new dynamically allocated array (newarr):

```
int ** tripleCols(int ** arr, int rows, int cols)
{
    int ** newarr = malloc(sizeof(int *) * rows * 3); // point a

    for(int r = 0; r < rows; ++r)
    {
        newarr[r] = malloc(sizeof(int) * cols * 3);
        newarr[r] = arr[r]; // point b
        free(arr[r]);
    }

    free(arr);
    return newarr;
}
```

The key errors in the code are at points a and b, noted in the comments.

- (a) (1 pt) What is the fix for the line of code for point a? (This one's simple, so no need to write out the new line of code, just describe the fix.)

Remove the “* 3” since the number of rows needs to be the same.

- (b) (2 pts) Why is the line of code for point b incorrect conceptually?

This line of code reassigns a pointer (newarr[r]) to point to a different location instead of copying over the values from arr to newarr.

- (c) (2 pts) Write two lines of code to replace this one line of code so that the function will work as planned.

```
for (int c=0; c<cols; c++)
    newarr[r][c] = arr[r][c];
```

Grading: part (a) is all or nothing, 1 pt, award 0, 1 or 2 pts for part (b) as you see fit, for part (c), give full credit if it's correct, 1 pt if there's some loop but it's not fully correct, or 0 if there's no loop.

2) (10 pts) DSN (Linked Lists)

We can store an integer in a linked list of nodes, where each node stores digit, in reverse order. For example, the integer 2163 would be stored in the linked list $3 \rightarrow 6 \rightarrow 1 \rightarrow 2$. Using the node struct shown below that is used to store numbers in this manner, write a **recursive** compareTo function that takes in pointers to two integer stored in this manner and returns a negative integer if the number in the list pointed to by num1 is less than the number in the list pointed to by num2, 0 if the two respective numbers are equal, or a positive integer if the number in the list pointed to by num1 is larger than the number in the list pointed to by num2. For example, compareTo($3 \rightarrow 6 \rightarrow 1 \rightarrow 2$, $4 \rightarrow 6 \rightarrow 1 \rightarrow 2$) should return a negative integer and compareTo($3 \rightarrow 6 \rightarrow 1 \rightarrow 2$, $9 \rightarrow 9 \rightarrow 9 \rightarrow 1$) should return a positive integer.

```
typedef struct node {
    int digit;
    struct node* next;
} node;

int compareTo(node* num1, node* num2) {

    if (num1 == NULL && num2 == NULL) return 0;

    if (num1 == NULL) return -1;

    if (num2 == NULL) return 1;

    int tmp = compareTo(num1->next, num2->next);

    if (tmp != 0) return tmp;

    return num1->digit - num2->digit;
}
```

Grading: 1 pt for each base case (order matters to avoid null ptr error)

3 pts for the recursive call

2 pts to return the recursive call answer when it's not 0.

2 pts to return appropriately when the least significant digit breaks the tie.

Note: The base cases can be done in a few ways.

Most students will probably use an if-else to compare num1->digit, num2->digit at the end there.

3) (10 pts) ALG (Stack)

Convert the following infix expression to postfix using a stack. Show the contents of the stack at the indicated points (A, B, and C) in the infix expression.

$$9 * 2 + 4 - 3 / (6 * (4 - 3)) + 1 * (7 - 4) / (5 + 6)$$
[illegible]

A

[illegible]

B

(
/
+

C

Note: A indicates the location in the expression **AFTER** the division operator and before the open parenthesis. B indicates the location in the expression **AFTER** the multiplication and before the open parenthesis. C indicates the location in the expression **AFTER** the open parenthesis and before the value 5.

Resulting postfix expression:

9	2	*	4	+	3	6	4	3	-	*	/	-	1	7	4	-	*	5	6	+	/	+
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

Note: There are exactly the correct number of boxes above. These should be filled with 12 numbers and 11 operators.

Grading: 1 pt for each stack, all or nothing, 7 pts total for the expression, take off 1 pt per error, cap at 0. (If you can fix the expression by moving 1 item to a different relative location, then that counts as one error.)

Computer Science Foundation Exam

May 17, 2025

Section B

ADVANCED DATA STRUCTURES

**NO books, notes, or calculators may be used,
and you must work entirely on your own.**

SOLUTION

Question #	Max Pts	Category	Score
1	10	ALG	
2	5	ALG	
3	10	DSN	
TOTAL	25	----	

You must do all 3 problems in this section of the exam.

Problems will be graded based on the completeness of the solution steps and not graded based on the answer alone. Credit cannot be given unless all work is shown and is readable. Be complete, yet concise, and above all be neat. For each coding question, assume that all of the necessary includes (stdlib, stdio, math, string) for that particular question have been made.

1) (10 pts) ALG (Binary Trees)

The following code, when passed the root of a binary tree and returns a result calculated by adding the values of some nodes and subtracting the values of others. If the function whatDoesItDo is called on the root of the binary tree shown below, for each value in the tree, indicate whether it gets added (A) or subtracted (S), **with respect to the original call on the root of the tree below**. No need to state the final return value. (**Grading Note:** +1 for each correct slot, +0 for each slot left blank, -1 for each incorrect slot, minimum score is 0.)

```
#include <stdio.h>
#include <stdlib.h>
typedef struct bintreenode {
    int data;
    struct bintreenode* left;
    struct bintreenode* right;
} bintreenode;


int whatDoesItDo(bintreenode* root) {

    if (root == NULL) return 0;
    if (root->left == NULL && root->right == NULL) return root->data;

    if (root->left == NULL) return root->data + whatDoesItDo(root->right);
    if (root->right == NULL) return root->data + whatDoesItDo(root->left);

    if (root->left->data > root->right->data)
        return root->data + whatDoesItDo(root->left) - whatDoesItDo(root->right);

    return root->data + whatDoesItDo(root->right) - whatDoesItDo(root->left);
}
```



For each open slot, either write the letter 'A' for added, or 'S' for subtracted.

30	<u>A</u>	16	<u>S</u>	22	<u>A</u>
19	<u>S</u>	18	<u>A</u>	40	<u>A</u>
25	<u>S</u>	6	<u>A</u>	14	<u>S</u>
23	<u>A</u> (Grading is described in the question +1 correct, 0 blank, -1 incorrect, cap at 0)				

2) (5 pts) ALG (Hash Tables)

Consider a hash table that uses the **quadratic probing technique** with the following hash function $f(x) = (3x+4)\%11$. (The hash table size is 11). If we insert the values 22, 11, 44, 32, 10, 21, and 33 into the table, in that order, show where these values would end up in the table.

Index	0	1	2	3	4	5	6	7	8	9	10
Value		32	10		22	11			44	33	21

Grading: 1 pt total for 22, 11 and 44

1 pt total for 32 and 10

1 pt for placement of 21

2 pts for placement of 33

To get the point, all the items in the list for that point have to be in the correct slot.

3) (10 pts) DSN (Tries)

We are maintaining a Trie for predicting the next letter(s) for a given string. The trie node struct and its properties are discussed below.


```
typedef struct trienode {
    int freq;
    int sum_prefix_freq;
    int cur_max_freq;
    struct trienode* next[26];
} trienode;
```

- **freq**: The frequency of the word represented by this node (i.e., how many times this specific word has been added to the dictionary). If this value is 0, it means that there is no word in the trie that ends at that node.
- **sum_prefix_freq**: The total frequency of all words in the dictionary that have this string as a prefix, including the string itself.
- **cur_max_freq**: The highest sum frequency among all child nodes of the current node.
- **next[26]**: These are typical children pointers of a trie node. This is an array of 26 pointers, each representing one of the possible next letters ('a' to 'z'). A pointer should be NULL if no words in the dictionary continue along that path. Typically, only a subset of these pointers will be active.

As an example, the following trie is constructed after inserting the following list of words and their frequency. The numbers inside the nodes are written in the sequence of `freq`, `sum_prefix_freq`, `cur_max_freq`.

List of words and their frequency added to the trie.

cap 15, cat 20, act 10, able 10, ace 2,



Your goal is to **complete** the recursive function on the next page that receives the root of a trie, `t`, a string, `str`, and an integer, `k`, the current position in the string, and returns the most likely letter that follows the input string. **You may assume a unique next letter appears the most number of times** (`cur_max_freq`). If there is no string in the trie that `str` is a proper prefix for, then return question mark character, "?".

For example, if the string passed to the function is:

- predict(root, "a", 0) should make a recursive call to predict(root->next[0], "a", 1), which should then return 'c' because 'c' is the most likely letter to follow "a" for the sample trie.
- predict(root, "ab", 0) will eventually return "l" after two recursive calls.
- predict(root, "ac", 0) will eventually return "t" after two recursive calls.
- predict(root, "ace", 0) will eventually return "?" after three recursive calls because "ace" is not a proper prefix of any word in the trie.
- predict(root, "ap", 0) will make one recursive call and then from there return "?", since "ap" isn't a prefix of any word in the trie. (In code, this case is slightly different than the previous one.)

```
char predict(trienode *t, char *str, int k) {

    if (t == NULL) return '?';

    if (k == strlen(str)) {

        // Checks if there is no string with this prefix.
        // Grading: 2 pts
        if ( t->cur max freq == 0 )
            return '?';

        // Looks through all possible next letters until it finds the one
        // has the most words that start with that prefix.
        // Note: part before the && is for short-circuiting to avoid null ptr.
        for (int i=0; i<26; i++) {

            // Grading: 2 pts
            if ( t->next[i] != NULL &&

                // Grading: 3 pts total
                t->next[i]->sum prefix freq == t->cur max freq )

                return (char)('a'+i);

        }

    }

    // We require a recursive call in this case.
    // Grading: 3 pts total
    return predict(t->next[str[k]-'a'], str, k+1);
}
```

Computer Science Foundation Exam

May 17, 2025

Section C

ALGORITHM ANALYSIS

**NO books, notes, or calculators may be used,
and you must work entirely on your own.**

SOLUTION

Question #	Max Pts	Category	Score
1	10	ANL	
2	5	ANL	
3	10	ANL	
TOTAL	25	----	

You must do all 3 problems in this section of the exam.

Problems will be graded based on the completeness of the solution steps and not graded based on the answer alone. Credit cannot be given unless all work is shown and is readable. Be complete, yet concise, and above all be neat. For each coding question, assume that all of the necessary includes (stdlib, stdio, math, string) for that particular question have been made.

1) (10 pts) ANL (Algorithm Analysis)

Consider a n -bit binary counter, which starts with the binary representation of 0 and increments by 1 until it reaches the binary value of $2^n - 1$. For $n = 3$, the counter would start at 000, and then change as follows:

$$000 \rightarrow 00\underline{1} \rightarrow 0\underline{1}0 \rightarrow 0\underline{1}1 \rightarrow \underline{1}00 \rightarrow 1\underline{0}1 \rightarrow 1\underline{1}0 \rightarrow 1\underline{1}1.$$

The underlined bits represent the ones that had to be changed. In particular, for this example, $1 + 2 + 1 + 3 + 1 + 2 + 1 = 11$ bits were changed as the counter progressed from 0 to $2^n - 1$. Let $f(n)$ equal the number of bits that are changed for an n -bit binary counter counting from 0 to $2^n - 1$. Find a closed-form formula for $f(n)$. (For example, something like $f(n) = 2^{n-1} + 2$. A formula in terms of n without any sort of recursive function definition.) Show all of your work and put a box around your final answer.

They key observation is that the lowest-order zero bit in the binary number controls the number of bits that get changed with the counter increments. Thus, when the binary counter is 1010110111, for example, and we note that the lowest-order zero bit is in the fourth position, counting from the right (1-based counting), we know that exactly 4 bits will flip when the counter increments to 1010111000.

In terms of n , the counter ends in a 0 2^{n-1} times (half of the 2^n numbers displayed on the counter). In these cases, 1 bit gets flipped.

In terms of n , the counter ends in 01 2^{n-2} times and will get flipped twice in these cases.

In terms of n , the counter ends in 011 2^{n-3} times and will get flipped three times in these cases.

This pattern persists for each value of k , as k ranges from 1 to n . (The counter ends in 011..1 exactly 2^{n-n} , or 1 time and all n bits flip this one time.)

It follows that $f(n) = \sum_{k=1}^n k 2^{n-k}$. Let's evaluate this sum. We first write it down, and then we take the whole expression and divide it by 2 and write down the corresponding sum below the original. Then we subtract the bottom equation from the top:

$$\begin{array}{rcl} f(n) = & 1 \times 2^{n-1} + 2 \times 2^{n-2} + 3 \times 2^{n-3} + \dots & + n \times 2^0 \\ f(n)/2 = & 1 \times 2^{n-2} + 2 \times 2^{n-3} + \dots & + (n-1) \times 2^0 + n \times 2^{-1} \\ \hline f(n) - f(n)/2 = & 2^{n-1} + 2^{n-2} + 2^{n-3} + \dots & + 2^0 - n/2 \\ f(n)/2 = & 2^{n-1} + 2^{n-2} + 2^{n-3} + \dots & + 2^0 - n/2 \end{array}$$

Use the geometric sum formula to evaluate the sum on the right except for the last term:

$$f(n)/2 = (2^n - 1) - n/2$$

Multiply by 2 to get

$$f(n) = 2(2^n - 1) - n = 2^{n+1} - 2 - n = \underline{\underline{2^{n+1} - n - 2.}}$$

Alternate Solution to #1

Instead of summing each number in the original problem individually ($1 + 2 + 1 + 3 + 1 + 2 + 1$, which we decomposed into $4 \times 1 + 2 \times 2 + 1 \times 3$), we can view the problem differently and count how many times each bit gets flipped. The least significant bit gets flipped every time, or $2^n - 1$ times. The second least significant bit gets flipped slightly less than half of that, exactly $2^{n-1} - 1$ times. (To see this, note that we flip this bit every other time, with the flip occurring on the second of each pair. Formally, we flip this bit $\left\lfloor \frac{2^n - 1}{2} \right\rfloor = 2^{n-1} - 1$ times. More generally, the k th least significant bit gets flipped exactly $\left\lfloor \frac{2^n - 1}{2^{k-1}} \right\rfloor = 2^{n-k+1} - 1$ times. Thus, we can add up the total number of bit flips by adding the number of times each individual bit itself gets flipped, giving us the following summation to evaluate:

$$\begin{aligned}
 & \sum_{k=1}^n 2^{n-k+1} - 1 \\
 & \sum_{k=1}^n (2^{n-k+1} - 1) = \left(\sum_{k=1}^n 2^{n-k+1} \right) - \left(\sum_{k=1}^n 1 \right) \\
 & = \left(\sum_{k=1}^n 2^k \right) - n \\
 & = 2(2^n - 1) - n \\
 & = 2^{n+1} - 2 - n \\
 & = 2^{n+1} - n - 2
 \end{aligned}$$

Yet a third way to view this problem is to let $T(n)$ be the answer to the question for an n -bit counter. Using the observation above where we note that the least significant bit flips every time ($2^n - 1$), notice that the remaining $n - 1$ bits are essentially playing the role of an $n - 1$ bit-counter. (Basically, the $n - 1$ most significant bits stay frozen every other step and then just count regularly on the even numbered steps. This means that $T(n) = T(n - 1) + 2^n - 1$. The solution to this recurrence is the summation above.

Grading: 4 pts for setting up summation or recurrence relation which corresponds to the answer to the question.

6 pts for evaluating the derived summation or recurrence relation.

Give partial credit as you see fit for both parts.

If initial summation is incorrect but that incorrect sum is evaluated correctly, give a maximum of four points out of six for evaluating the summation, depending on complexity of it.

2) (5 pts) ANL (Algorithm Analysis)

A $O(\sqrt{n})$ search algorithm took 45 milliseconds to complete a search amongst $n = 4 \times 10^6$ entries. How long would it be expected for this algorithm execute a search amongst a database of 10^8 entries, in milliseconds?

Let $T(n) = c\sqrt{n}$ be the amount of time the algorithm takes to execute on a input of size n . Using the given information we have:

$$T(4 \times 10^6) = c\sqrt{4 \times 10^6} = 45 \text{ ms}$$

$$c\sqrt{4 \times 10^6} = 45 \text{ ms}$$

$$(2 \times 10^3)c = 45 \text{ ms}$$

$$c = \frac{45}{2000} \text{ ms}$$

Now, we solve for $T(10^8)$:

$$T(10^8) = \frac{45}{2000} \text{ ms} \sqrt{10^8} = \frac{45 \times 10^4}{2 \times 10^3} \text{ ms} = 45 \times 5 \text{ ms} = \mathbf{225 \text{ ms}}$$

Grading: 1 pt set up equation for c.

1 pt solve for c.

1 pt plug in $n = 10^8$

2 pts to get to correct final answer simplified as 225 ms. (1 pt for intermediate form)

3) (10 pts) ANL (Recurrence Relations)

Determine the following summation in terms of n , **in factorized form**. (Do NOT multiply the answer out into polynomial form. Note: Your answer should NOT have a fraction in it.)

$$\sum_{i=1}^{2n-1} (i + 3i^2)$$

$$\begin{aligned} \sum_{i=1}^{2n-1} (i + 3i^2) &= \left(\sum_{i=1}^{2n-1} i \right) + \left(\sum_{i=1}^{2n-1} 3i^2 \right) \\ &= \frac{(2n-1)(2n)}{2} + \frac{3(2n-1)(2n)(2(2n-1)+1)}{6} \\ &= n(2n-1) + \frac{3(2n-1)(2n)(4n-2+1)}{6} \\ &= n(2n-1) + (2n-1)(n)(4n-1) \\ &= n(2n-1)(1+4n-1) \\ &= n(2n-1)(4n) \\ &= 4n^2(2n-1) \end{aligned}$$

Grading: 1 pt split sum

2 pts formula sum of i

2 pts formula sum of i^2

2 pts to get to non-fractional form (canceling 2, 6)

2 pts factor out $n(2n-1)$

1 pt to simplify to final form

Note: Grade was 7 pts out of 10 for polynomial form.

Computer Science Foundation Exam

May 17, 2025

Section D

ALGORITHMS

**NO books, notes, or calculators may be used,
and you must work entirely on your own.**

SOLUTION

Question #	Max Pts	Category	Score
1	10	DSN	
2	10	ALG	
3	5	ALG	
TOTAL	25	----	

You must do all 3 problems in this section of the exam.

Problems will be graded based on the completeness of the solution steps and not graded based on the answer alone. Credit cannot be given unless all work is shown and is readable. Be complete, yet concise, and above all be neat. For each coding question, assume that all of the necessary includes (stdlib, stdio, math, string) for that particular question have been made.

1) (5 pts) DSN (Recursive Coding)

Write a **recursive function** below so that it returns the maximum integer k such that $\text{base}^k \leq \text{ans}$. For example if base is 3 and ans is 123, then 4 should be returned since $3^4 = 81$ and $3^5 = 243$. The restricted bounds below are to ensure that integer overflow isn't an issue.

```
// Pre-condition: 1 < base <= 1000, 1 <= ans <= 1000000
int maxpowLTE(int base, int ans) {

    // Grading: 1 pt if, 1 pt return.
    if (ans < base)
        return 0;

    // Grading: 3 pts total, 1 pt return 1 +, 1 pt rec call,
    // 1 pt parameters to rec call.
    return 1 + maxpowLTE(base, ans/base);

}
```

2) (10 pts) DSN (Sorting)

Consider the problem of sorting a competition struct. A competition struct has three integer components: **probSolved**, **totalTime** and **difficulty**. One struct is greater than another if it has more problems solved (**probSolved**) than another. If the problems solved is the same between two structs and the total time is less, then that struct is greater than the other. Finally, between two structs with the same number of problems solved and total time, if one has greater difficulty, it is greater than the other struct. If all three components are equal, the structs are equal and neither is greater than the other. Write a function called `greaterThan`, which takes in pointers to two competition structs and returns 1 if the struct pointed to by `ptrA` is greater than the struct pointed to by `ptrB`, according to these rules, and 0 otherwise. For example, (pS=3, tT = 200, d = 8) is greater than (3, 200, 7) but is NOT greater than (3, 199, 7). Rather (3, 199, 7) is greater than (3, 200, 8).

```
typedef struct competition {
    int probSolved;
    int totalTime;
    int difficulty;
} competition;

int greaterThan(competition* ptrA, competition* ptrB) {

    // Grading: 1 pt if, 1 pt return.
    if (ptrA->probSolved > ptrB->probSolved) return 1;

    // Grading: 1 pt if, 1 pt return.
    if (ptrA->probSolved < ptrB->probSolved) return 0;

    // Grading: 1 pt if, 1 pt return.
    if (ptrA->totalTime < ptrB->totalTime) return 1;

    // Grading: 1 pt if, 1 pt return.
    if (ptrA->totalTime > ptrB->totalTime) return 0;

    // Grading: 2 pts total
    if (ptrA->difficulty > ptrB->difficulty) return 1;
    return 0;
}
```

Grading Note: The order of these statements matters (some), if each statement needed is present, but the order makes the code incorrect, then take off an appropriate number of points for the incorrect order (say 1, 2 or 3 pts depending on the severity.)

3) (10 pts) DSN (Bitwise Operators)

User IDs are stored in a system as single integers in between 0 and $2^{25} - 1$, inclusive. Thus, each User ID can be viewed as a bitstring of length 25 (using an integer variables 25 least significant bits). When a new user ID is added, in order not to cause confusion, it's required that it differs in **at least three bits** compared to all other user IDs in the system. Write a function that takes in an array of current user IDs (**curIDs**), the length of that array (**n**), and a potential ID to be added (**pid**) and returns 1 if the potential ID can be added to the current list based on the previously given criteria, OR 0 if it can't be added if there exists a current ID with which the new ID differs in 2 or fewer bits. (For example, **10010000** differs with **00110000** in exactly 2 positions: 2^5 and 2^7 . Thus, if the former was in the current ID list, the latter would NOT be an allowable password to add. **PLEASE DO NOT WRITE ANY AUXILIARY FUNCTIONS.**

```
int canBeAdded(int* curIDs, int n, int pid) {

    for (int i=0; i<n; i++) {          // Grading: 1 pt

        int XOR = curIDs[i]^pid;      // Grading: 1 pt

        int diff = 0;                  // Grading: 1 pt
        for (int j=0; j<25; j++)      // Grading: 1 pt
            if (XOR & (1<<j))          // Grading: 2 pts
                diff++;                // Grading: 1 pt

        if (diff < 3) return 0;        // Grading: 2 pts
    }

    return 1;                          // Grading: 1 pt

}
```

Grading Note: If a line of code is not in the proper nesting of loops, take off 1 or 2 points as you see fit.