

ZJU-UIUC Institute

Second Midterm Exam, ECE 220

Thursday 3 December 2020

Name (pinyin and Hanzi):

SOLUTION IS IN RED

Student ID:

Lab TA Name:

- Be sure that your exam booklet has 12 pages.
- Write your name, student ID, and lab section TA name on the first page.
- Do not tear the exam apart other than to remove the reference sheet.
- This is a closed book exam. You may not use a calculator.
- You are allowed TWO handwritten A4 sheets of notes (both sides).
- **YOU MAY NOT USE EXTRA PAPER! WRITE ON THE EXAM!**
- Absolutely no interaction between students is allowed.
- Show all work, and clearly indicate any assumptions that you make.
- Challenge problems are marked with ***.
- Don't panic, and good luck!

Problem 1	19 points	_____
Problem 2	25 points	_____
Problem 3	25 points	_____
Problem 4	30 points	_____
Correct Room	1 point	_____

Total	100 points	_____
-------	------------	-------

Problem 1 (19 points): Short Answer Questions

Tingkai is working on a filesystem for his operating system in ECE391. His plan is to organize the filesystem as a set of data blocks, each containing 4 kB (4,096 bytes), using the following structure:

```
struct data_block_t {
    uint8_t bytes[0x1000];
};
```

1. (3 points) After struggling for a while, Tingkai found a function that returns a pointer to the first block of the file system:

```
struct data_block_t* fblock = get_fs_pointer ();
```

Tingkai then wrote following variable declaration below to copy a block from the filesystem:

```
struct data_block_t blk = *(fblock + 4*N + 2);
```

His code doesn't look elegant. Please help him by rewriting the declaration using array notation (**N** is a variable).

```
struct data_block_t blk = fblock[4 * N + 2] ;
```

Now consider an array of data blocks:

```
struct data_block_t my_blocks[4];
```

Assuming a machine with byte-addressable memory (one memory address in memory holds one byte), the array appears as shown to the right.

2. (3 points) What is the value of

```
my_blocks[2].bytes[1]?
```

Write "bits" if the value cannot be determined.

0xBC

my_blocks[0]	→	0x12
		0x34
		...
my_blocks[1]	→	0x56
		0x78
		...
my_blocks[2]	→	0x9A
		0xBC
		...
my_blocks[3]	→	0xEE
		0xFF
		...

3. (3 points) Given the declaration:

```
struct data_block_t* foo = my_blocks + 3;
```

and assuming that **my_blocks** has value **0xECE220**, what is the value stored in **foo** (in hexadecimal)? Show your work for partial credit.

0xECE220 + 0x1000 * 3 = 0xED1220

Problem 1, continued:

4. (4 points) In ECE220, Bob learned that a one-dimensional array can be passed as a pointer to a function. He wanted to extend this idea to two-dimensional arrays and wrote the following code:

```
void foo (char** arr) {
    // Some code
}

int main () {
    char a[10][20];
    foo (a);
    return 0;
}
```

Is this code correct? USING 30 WORDS OR FEWER, explain.

No. a is a pointer to an array of arrays of 20 characters (contiguous in memory), not a
pointer to a pointer

5. (6 points) Bob also struggled with the idea of dynamic allocation. On a recent midterm, he was asked to write a function to free a linked list of `thing_ts`, given a pointer to a variable holding a pointer to the head of the list. The function should also set the original list head variable to NULL.

Bob wrote the code shown below. Unfortunately, it contains TWO BUGS.

USING 20 WORDS OR FEWER (per bug), explain each bug and how to fix the problem.

Bug 1: uses thing->next in loop update after freeing it in loop body; add a temporary
variable

Bug 2: writes NULL to head_ptr parameter instead of *head_ptr; add an asterisk

```
typedef struct thing_t thing_t;
struct thing_t {
    // Other fields don't matter.
    thing_t* next;
};

void free_list (thing_t** head_ptr)
{
    for (thing_t* thing = *head_ptr; NULL != thing; thing = thing->next) {
        free (thing);
    }
    head_ptr = NULL;
}
```

Problem 2 (25 points): Arrays and Debugging with Deep Neural Networks

1. (15 points) In this problem, you must implement a convolution, an important tool for image processing and deep neural networks (DNNs). Given a $\text{size} \times \text{size}$ input matrix **in** and a 3×3 mask matrix **mask**, your function must calculate a $\text{size} \times \text{size}$ matrix **out**, as described below.

In a convolution, each output element is calculated by first aligning the center of the mask matrix over the corresponding input element (the element with the same row and column indices as the output element being calculated), as illustrated by the shaded region of the input in the example below for $(\text{row}, \text{col}) = (0, 0)$. Notice that part of the mask may fall outside of the input for some output elements (including the example shown), and that elements outside the actual input matrix are treated as 0. After alignment, each element of the mask is multiplied by the corresponding input element (or 0), and all nine products are summed to produce the single output element.

0	0	0						
0	1	1	1					
0	1	2	1					
	1	1	1					

*

1	1	2
-2	1	1
2	1	1

=

5	5	4
8	9	2
7	5	2

in
mask
out

Specifically, in the figure above, all three matrices—**in**, **mask**, and **out**—are 3×3 . To calculate the $(0, 0)$ element of the output matrix **out** (the shaded element), we align the center of the mask over the $(0, 0)$ element of the input matrix **in**—the shaded region shows the position of the **mask** matrix.

Elements of the shaded region that fall outside of the input matrix use the value 0 instead of input matrix values, as shown in the figure. To compute the output value, we multiply each of the values from the mask by the corresponding element of the input (or 0), then sum up the nine products. In this case, starting from the upper left, we obtain $0 \times 1 + 0 \times 1 + 0 \times 2 = 0$ from the first row, $0 \times (-2) + 1 \times 1 + 1 \times 1 = 2$ from the second row, and $0 \times 2 + 1 \times 1 + 2 \times 1 = 3$ from the third row, for a total of 5, which we write into **out** at position $(0, 0)$.

The problem is on the next page.

Problem 2, continued:

0	0	0	
0	1	1	1
0	1	2	1
	1	1	1

*

1	1	2
-2	1	1
2	1	1

=

5	5	4
8	9	2
7	5	2

in
mask
out

(Image replicated for your convenience.)

USING AT MOST SIX LINES (not counting braces and variable declarations), which must fit within the blanks provided, complete the inner loop body below to handle one product term of each output element. Include comments for more partial credit.

Each of the three matrices—**in**, **mask**, and **out**—is given as a one-dimensional array. Calculate array indices in the same manner as discussed in class and used in your MPs.

You may ignore the possibility of overflow in your calculations.

```
void conv_layer (int32_t* in, int32_t* mask, int32_t* out, int32_t size)
{
    int32_t x, y, p, q;

    // For each value in the input matrix
    for (y = 0; y < size; y++) {
        for (x = 0; x < size; x++) {

            // Initialize out to 0
            out[y * size + x] = 0;

            // For each value in the mask matrix
            for (p = 0; p < 3; p++) {
                for (q = 0; q < 3; q++) {

                    if (0 <= y + p - 1 && size > y + p - 1 &&
                        0 <= x + q - 1 && size > x + q - 1) {

                        out[y * size + x] += mask[p * 3 + q] *
                            in[(y + p - 1) * size + (x + q - 1)];

                    }

                }
            }
        }
    }
}
```

Problem 2, continued:

2. (10 points) Pooling is a second important operation in DNNs. In max pooling, a square submatrix of values is replaced with a single value—the **maximum** among the values in the square submatrix. The picture below, for example, shows 2×2 max pooling applied to a 5×5 input matrix.

29	35	26	167	55
0	100	45	22	33
14	14	7	34	21
14	14	22	56	22
16	-6	2	155	23

in

-->

100	167
14	56

out

Your friend has implemented 2×2 max pooling in the subroutine below. Given a `size × size` input matrix `in`, the subroutine produces an appropriately sized output matrix `out` (if `size` is odd, the code should ignore the last row and column).

Unfortunately, **your friend's subroutine HAS TWO BUGS**. For each bug, give one example of an input matrix `in` (for example: `in = {1, 2, 3, 4}`) that exposes the bug. Then, **USING TWENTY WORDS OR FEWER**, explain the bug.

Bug 1: `in = {` 1 `}`

Reason 1: Rounds out_size up instead of down, producing 1 rather than 0 when size is 1.

Bug 2: `in = {` -1, -1, -1, -1 `}`

Reason 2: Initializes val_max to 0 instead of minimum int32_t value, so calculated max is wrong.

```
void maxpooling_layer(int32_t *in, int32_t *out, int32_t size) {
    int32_t out_size = (size + 1) / 2;
    int32_t x, y, p, q;
    int32_t val_max, val_cmp;

    // For each value in the output matrix
    for (x = 0; x < out_size; x++) {
        for (y = 0; y < out_size; y++) {
            val_max = 0;

            for (p = 0; p < 2; p++) {
                for (q = 0; q < 2; q++) {
                    // value to be compared
                    val_cmp = in[(2 * y + p) * size + (2 * x + q)];
                    val_max = (val_max > val_cmp ? val_max : val_cmp);
                }
            }
            out[y * out_size + x] = val_max;
        }
    }
}
```

Problem 3 (25 points): Functions and Dynamic Resizing

Recall the AI quantum magic button that you implemented during Midterm 1 for D-331. To further improve security, you decide to allocate a unique ID (a `uint32_t`) for each student, and to store these IDs in a dynamically resized array.

The ID list is stored in a dynamically resized array using the following file-scope variables.

```
uint32_t* id_list;    // pointer to the dynamically allocated ID array
uint32_t num_ids;    // current number of elements stored in ID array
uint32_t max_ids;    // current ID array size
```

Complete each function below by filling in the blanks as necessary. Not all blanks may be needed. You may **USE ONLY THE BLANKS PROVIDED**. Additional code will earn no credit.

1. **(6 points)** Complete the function below to check whether the ID given by parameter `id` is already present in the array of IDs. Return 1 if it is present, or 0 if it is not present.

```
int32_t is_duplicate (uint32_t id) {
    uint32_t i;

    for (i = 0; num_ids > i; i++) {
        if (id == id_list[i]) {
            return 1;
        }
    }

    return 0;
}
```

2. **(4 points)** Complete the function below to add the ID given by parameter `id` to the end of the array of IDs. You may assume that the array has been allocated and contains enough space for the new ID. Update file-scope variables as necessary.

```
void insert_id (uint32_t id) {
    id_list[num_ids++] = id;
}
```

Problem 3, continued:

```
uint32_t* id_list;    // pointer to the dynamically allocated ID array
uint32_t  num_ids;    // current number of elements stored in ID array
uint32_t  max_ids;    // current ID array size
```

(List of file-scope variables replicated for your convenience.)

3. **(15 points)** Finally, complete the function below to insert the ID given by the parameter `id` to the end of the array of IDs. The function should return 1 for success, and 0 for failure.
- Duplicate IDs should be rejected (by returning failure).
 - You may assume that the array ID pointer is valid when your function is called (it will not be NULL).
 - If the array does not have enough space for a new ID, grow the array by a factor of 3, updating file-scope variables as necessary. Return failure if no memory is available.
 - See the reference sheet for dynamic allocation functions available in the C standard library.

Your code **MUST USE** the helper functions that you developed for **Part 1** and **Part 2**.

```
int32_t register_id (uint32_t id) {

    _____

    _____

    // Check for duplicate IDs.

    if (is_duplicate (id)) {

        return 0;

    }

    // Resize the array if necessary.

    if ( max_ids == num_ids ) {

        uint32_t* copy = realloc (id_list, 3 * max_ids * sizeof (*id_list));

        if (NULL == copy) {

            return 0;

        }

        id_list = copy;

        max_ids *= 3;

    }

    _____

    // Insert the ID.

    insert_id (id);

    return 1;

}
```


Problem 4 (30 points): Merge Sort on Linked Lists

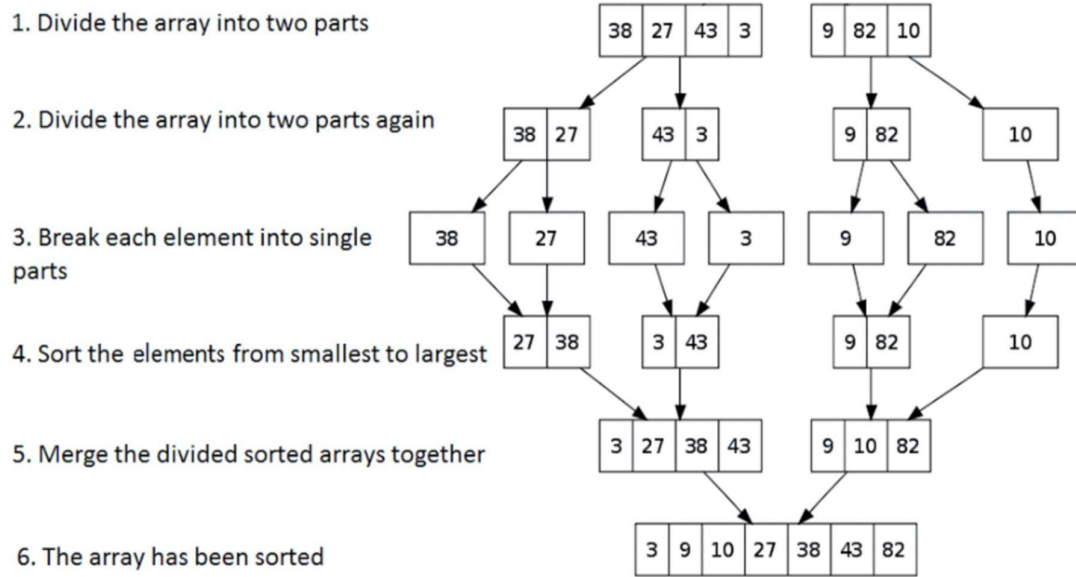
In this problem, you must implement a recursive merge sort on a linked list of values. Merge sort works as follows:

- Divide the original data into two unsorted parts (of roughly equal size).
- Sort both parts (recursively).
- Merge the two sorted parts into a sorted whole.

Note that any list consisting of a single element is already sorted, and forms a stopping condition for the recursion.

For each part of this problem, complete each function below by filling in the blanks as necessary. Not all blanks may be needed. You may **USE ONLY THE BLANKS PROVIDED**. Additional code will earn no credit.

As a further aid to understanding, the diagram below illustrates the merge sort process for an array.



Your recursive merge sort will operate on a linked list of the structure shown below, each of which has an integral value and a next pointer.

```
typedef struct element_t element_t;
struct element_t {
    int32_t value;
    element_t* next;
};
```

Problem 4, continued:

```
typedef struct element_t element_t;
struct element_t {
    int32_t    value;
    element_t* next;
};
```

(Structure definition replicated for your convenience.)

Before implementing the main recursive merge sort function, you must implement two helper functions.

1. ***** (12 points)** Complete the function below to divide the list starting with **head** into two sublists. The sublists should have equal length if the original list has even length. If the original list has odd length, the extra element should be put into the first sublist. A pointer to the head of the first sublist should be written to the address given by **firstp**, and a pointer to the head of the second sublist should be written to the address given by **secondp**.

Not all blanks may be needed. You may **USE ONLY THE BLANKS PROVIDED**. Additional code will earn no credit.

For this function, the list given by **head** must not be empty.

For full credit, do not write any additional loops (other than the one given).

```
void divide_list (element_t* head, element_t** firstp, element_t** secondp)
{
    element_t* middle = head;
    element_t* end     = head->next;
    element_t* sublist;

    while (end != NULL) {

        end = end->next;

        if (NULL == end) {

            break;

        }

        end = end->next;

        middle = middle->next;
    }

    sublist      = middle->next;
    middle->next = NULL;           // middle node ends one sublist

    *firstp = head;

    *secondp = sublist;
}
```

Problem 4, continued:

```
typedef struct element_t element_t;
struct element_t {
    int32_t value;
    element_t* next;
};
```

(Structure definition replicated for your convenience.)

Not all blanks may be needed. You may **USE ONLY THE BLANKS PROVIDED**. Additional code will earn no credit.

- (8 points)** Complete the recursive function below to merge two sorted lists, **fst** and **sec**, into a single sorted list and return a pointer to the merged list. **All lists are sorted in increasing order of their value fields.**

```
element_t* merge_list (element_t* fst, element_t* sec)
{
    element_t* result;

    if (fst == NULL)

        return sec;
    if (sec == NULL)

        return fst;

    if (fst->value < sec->value) {

        result = fst;

        result->next = merge_list (fst->next, sec);
    } else {

        result = sec;

        result->next = merge_list (fst, sec->next);
    }
    return result;
}
```

- (10 points)** Now you are ready to write merge sort. Complete the recursive function below to sort the list **head**, which may be empty, using merge sort, and return a pointer to the sorted list. Your code **MUST USE** the helper functions that you wrote in **Part 1** and **Part 2**.

```
element_t* merge_sort (element_t* head)
{
    element_t* fst;
    element_t* sec;

    if (NULL == head || NULL == head->next) {
        return head;
    }
    divide_list (head, &fst, &sec);

    fst = merge_sort (fst);

    sec = merge_sort (sec);

    return merge_list (fst, sec);
}
```

dynamic allocation routines from C's standard I/O library

```
// returns pointer to new memory, or NULL on failure
void* malloc (size_t size);

// returns pointer to 0-filled new memory, or NULL on failure
void* calloc (size_t nmemb, size_t size);

// returns pointer to resized block, or NULL on failure
void* realloc (void* ptr, size_t size);

// frees previously allocated block
void free (void* ptr);
```