

**Third Midterm Exam**  
November 28, 2018  
CS162 Operating Systems

<b>Your name</b>	
<b>SID</b>	
<b>CS162 login (e.g., s162)</b>	
<b>TA Name</b>	
<b>Discussion section time</b>	

**This is a closed book and two 2-sided handwritten notes** examination. You have 80 minutes to answer as many questions as possible. The number in parentheses at the beginning of each question indicates the number of points for that question. You should read **all** of the questions before starting the exam, as some of the questions are substantially more time consuming.

Write all of your answers directly on this paper. ***Make your answers as concise as possible.*** If there is something in a question that you believe is open to interpretation, then please ask us about it!

**Good Luck!!**

<b>Question</b>	<b>Points assigned</b>	<b>Points obtained</b>
<b>P1</b>	<b>22</b>	
<b>P2</b>	<b>20</b>	
<b>P3</b>	<b>18</b>	
<b>P4</b>	<b>14</b>	
<b>P5</b>	<b>12</b>	
<b>P6</b>	<b>14</b>	
<b>Total</b>	<b>100</b>	

## Problem 1: True/False questions (22 points)

**MARK THE CHECKBOX NEXT TO YOUR ANSWER.** For each question: 1 point for true/false correct, 1 point for explanation. An explanation **cannot** exceed 2 sentences.

- a) The UNIX Fast File System puts i-node metadata in the same cylinder group as the files they describe.

**TRUE**

**FALSE**

Why?

- b) As we add elements to a struct list in Pintos, the struct list itself will also grow in size.

**TRUE**

**FALSE**

Why?

- c) In a queueing system it is possible to have unbounded service times even when the arrival rate is less than the service rate.

**TRUE**

**FALSE**

Why?

- d) According to the end-to-end argument the reliability for a file transfer application should be implemented only in the network.

**TRUE**

**FALSE**

Why?

- e) Congestion control ensures that the sender will not overflow the receiver's buffer.

**TRUE**

**FALSE**

Why?

f) Consider a resource allocation system in which we can always preempt a process holding a resource. Such a system can never reach a deadlock state.

**TRUE**

**FALSE**

Why?

g) Retrieving a data block using the file allocation table (FAT) may require more than two disk accesses.

**TRUE**

**FALSE**

Why?

h) The shortest remaining job first scheduling discipline can lead to starvation.

**TRUE**

**FALSE**

Why?

i) The elevator algorithm (for disk scheduling) in general leads to lower seek times than the FIFO algorithm.

**TRUE**

**FALSE**

Why?

j) For a single processor, write-back caching results in fewer writes to disk than write-through caching.

**TRUE**

**FALSE**

Why?

k) RAID5 with one parity block can recover from one disk failure.

**TRUE**

**FALSE**

Why?

## Problem 2: Disk & RAID (20 points)

Several CS162 students decided they are going to start a start-up of streaming videos called Calfix. This start-up will help stream videos of CS162 lectures around the world. To do so, they need to understand how to design their storage system. Since these are CS162 students, they don't have too many resources, so they decided to start with very old magnetic disks for their storage system.

They bought disks with the following specs:

- 15ms controller delay
- 10ms seek time
- 3000 rpm rotation speed
- 100 megabyte/s transfer rate
- 200 kilobyte sector size.

a) What is the *average* rotation time (in ms) of the disk to read 1 sector? (2 point)

b) What is the transfer time to read 1 sector? (2 point)

c) Assuming a lecture is stored on 20 contiguous sectors on the same track, what is the average total time to read a movie? (2 points)

The Calfix team was able to fit the lecture of the last 20 years of CS162 lectures into their “data center” on 20 disks. Also, our Calfix founders remembered they learned in CS162 about something called RAID. They started with RAID0 which stripes a lecture on all 20 disks, i.e., each disk gets an equal number of sectors of each lecture. However, there has been a bug in the implementation of RAID0 and now each sector is stored randomly (i.e., non-contiguously) on each disk.

- d) What is the worst-case time it takes to read a lecture? Assume each lecture is stored on 20 sectors, like in question 3 (4 points)

- e) The Calfix engineers decided to increase the quality of the movies, and now it takes 40 sectors to store each movie. What is the worst case time takes to read a higher quality movie movie? (2 points)

The Calfix engineers would like to make their storage fault-tolerant so they have decided to implement RAID5 by using a parity sector for every 19 data sectors. Again the sectors of a movie are striped across all disks, and we are still incurring the previous bug in which the sectors are randomly placed on each disk. Assume each lecture is stored on 40 sectors.

- f) Assume a server fails. What is the time it takes to read a movie? (3 points)

- g) Finally the Calfix engineers have found and fixed the bug, so now all sectors of a movie that are stored on the same disk are stored contiguously. What is the *maximum* time it takes to read a movie in this case? (3 points)

## Problem 3: Malloc (18 points)

Recall the simple first-fit memory allocator you implemented in Homework 3.

- a) (4 points) Given the following series of allocations, write out a series of calls to `mm_free` and `mm_malloc` that would cause your memory allocator to request more space using `sbrk` even though enough memory has already been requested. For this question, you can assume the size of each block's header is 32 bytes. Also, briefly explain in 1 or 2 sentences why this happens.

```
void *p1 = mm_malloc(1000);  
void *p2 = mm_malloc(1000);  
void *p3 = mm_malloc(1000);
```

---

---

---

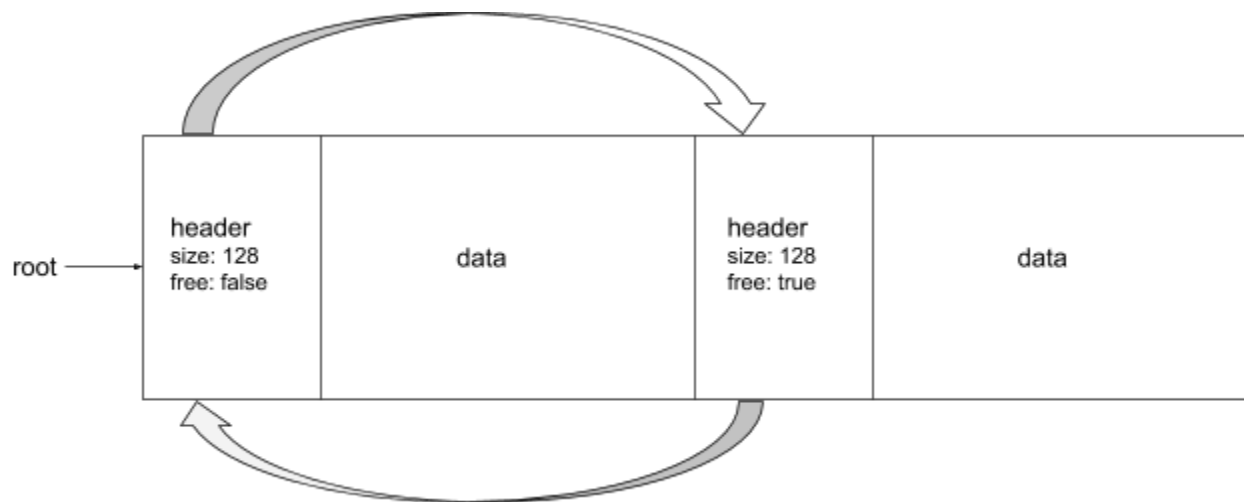
---

- b) (2 points) You want to fix this fragmentation issue because it can potentially waste a lot of memory if it goes unchecked for too long, but you don't want to write a new allocation algorithm. Instead, you change your memory allocator so that around every one hundred calls to `mm_malloc`, a compaction routine is called that causes all of the allocated blocks to be moved to the beginning of the heap with no free blocks in between each non-free block. Briefly (1 sentence) explain what could go wrong with this strategy. Bad performance will not be considered a potential issue.

- c) Describe the layout of the heap after the given calls to `mm_malloc` and `mm_free` using the first-fit allocator with free block coalescing specified in homework 3 by indicating the size of each block and whether each block is free or not in the tables below. Consider each part independently from the rest, and assume that metadata headers are 32 bytes large. If there are fewer blocks in the heap than there are columns, leave the extra columns blank. A block with a lower number corresponds to a block that comes earlier in the linked list. **Block numbers may or may not correspond to pointer numbers.** An example of an acceptable answer as well as a drawing of the corresponding layout on the heap for the following allocations is given below:

```
void *p1 = mm_malloc(128);
void *p2 = mm_malloc(128);
mm_free(p2);
```

	Block 1	Block 2	Block 3	Block 4
Size	128	128		
Free	false	true		



a) (3 points)

```
void *p1 = mm_malloc(128);  
void *p2 = mm_malloc(128);  
void *p3 = mm_malloc(128);  
void *p4 = mm_malloc(128);  
mm_free(p2);  
mm_free(p3);
```

	Block 1	Block 2	Block 3	Block 4
Size				
Free				

b) (3 points)

```
void *p1 = mm_malloc(128);  
void *p2 = mm_malloc(128);  
void *p3 = mm_malloc(128);  
mm_free(p2);  
void *p4 = mm_malloc(32);
```

	Block 1	Block 2	Block 3	Block 4
Size				
Free				

c) (3 points)

```
void *p1 = mm_malloc(128);  
void *p2 = mm_malloc(128);  
void *p3 = mm_malloc(128);  
mm_free(p2);  
void *p4 = mm_malloc(256);
```

	Block 1	Block 2	Block 3	Block 4
Size				
Free				



d) (3 points)

```
void *p1 = mm_malloc(128);  
void *p2 = mm_malloc(128);  
void *p3 = mm_malloc(128);  
void *p4 = mm_malloc(128);  
mm_free(p2);  
mm_free(p3);  
void *p5 = mm_malloc(255);
```

	Block 1	Block 2	Block 3	Block 4
Size				
Free				

## Problem 4. File System (14 points)

Donald wants to make some money by what he has learned from CS162. He has a customer named Hillary. Donald decided to use a new SSD drive to implement the file system. The SSD has a block size of 1KB, and the latencies to read and write a block are 50 usec, and 100 usec, respectively.

Notes:

- Ignore the block transfer time.
  - All throughput results should be given in MB/sec.
  - Recall that 1 sec = 1,000,000 usec.
- a. (3 points) Hillary runs a program that reads 64B *random* data chunks, i.e., each chunk is stored in a different block. What is the read throughput of Hillary's program?

- b. (2 points) Hillary is not happy with the throughput. To improve throughput, Donald helps Hillary to redesigning her application to store the chunks contiguously in a block. What is Hillary's program's throughput after this change?

To implement the file system, Donald decides on the following inode implementation. Recall that the block size is 1 KB.


```
struct inode_disk {
    off_t length; /* File size in bytes. */
    block_sector_t direct[15]; /* 12 direct pointers */
    block_sector_t indirect; /* a singly indirect pointer */
    block_sector_t doubly-indirect /* a doubly indirect pointer */
    block_sector_t triply-indirect /* a triply indirect pointer */
    uint32_t unused[114]; /* Not used. */
};
```

- c. (4 points) What is the maximum file size that this file system can support? (We will accept unsimplified answers.)

- d. (5 points) How long does it take to write an 1MB file?

## Problem 5: Queuing Theory (12 points)

Consider a single queue/single server system as below. The mean service time of each request is 20ms.



- a. (2 points) Assume the mean arrival rate is  $\lambda = 200/3$  request/sec. As the time goes to infinity, what is the queueing delay?

First engineer improves the server performance so that service time is brought down to 10ms. Assume that both arrival and service times are *fixed* (i.e., a request arrives every 15ms and it takes exactly 10ms to serve a request), and that at time 0 there are 3 requests in the queue.

- b. (2 points) What's the queueing delay when the system stabilizes (i.e., when queueing delay becomes a constant)?

- c. (4 points) How long does it take to stabilize?

Over time the load is changing and the engineers observe that requests spend now 30ms in the system (this includes both the time in the queue and at the server) on average. Assuming the mean of the arrival rate has not changed.

- d. (4 points) How many requests are in the system on average?


## Problem 6: TCP Flow Control (14 points)

Assume we use the TCP flow control protocol to transfer 200 bytes. Next, assume the followings parameters:


- Packet size: 50 bytes
- Maximum receiver buffer (MaxRcvBuf) is 150 bytes.
- The receiving application reads from the receiver buffer 100 bytes at a time. (Recall that the receiving application reads only bytes which are insequence, i.e., with no gaps).
- Acknowledgement contains the sequence # of the next expected in-sequence byte.
- Sender can send a new packet every 1ms, and the delay between the sender and receiver is 2m each way, so it takes 4ms for the sender to get an acknowledgement.
- Each tick on the timing diagram is 1 ms.
- The sender retransmit a packet in one of the following two cases:
  - If it doesn't receive an ack for that packet in 5ms.
  - If it gets a duplicate acknowledgement, i.e., an acknowledgment with the same sequence number as a previous one.

a) (4 points) Use the following figure to draw the packet diagram of the transfer. The first packet transfer and its acknowledgment is shown for you. Assume no losses.


- i) Note: *LastByteSent*, *LastByteAcked*, *LastByteWritten*, *NextByteRcvd*, *LastByteRead*, and *LastByteExpected* are shown for your convenience. You do not need to update them after each data packet / ack is sent /received in your solution.



b) (5 points) Draw the packet transfer diagram assuming the *last data packet* is lost.



- c) (5 points) Draw the packet transfer diagram assuming the first packet is lost. Assume the receiving process consumes 100 bytes of in order received data before the ack is sent out.





nO mOrE AnIMe



