

Fall 2016

Anthony D. Joseph

Midterm Exam #2 Solutions

October 25, 2016
CS162 Operating Systems

Your Name:	
SID AND 162 Login:	
TA Name:	
Discussion Section Time:	

General Information:

This is a **closed book and one 2-sided handwritten note** 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!!

QUESTION	POINTS ASSIGNED	POINTS OBTAINED
1	27	
2	30	
3	25	
4	18	
TOTAL	100	

Solutions NAME: _____

1. (27 points total) True/False and Why?

a. (12 points) True/False and Why? **CIRCLE YOUR ANSWER.**

i) It is possible for a FCFS scheduler to achieve a lower average turnaround time (the time a process takes to complete after it arrives) than a Round Robin scheduler, even if you assume there is no context switching overhead.

TRUE

FALSE

Why?

TRUE. Three examples: cache is shared with Round Robin, not FCFS; the jobs could arrive in order shortest to largest, reducing FCFS to SRTF; jobs could all be the same length, with that length being longer than the RR quanta. The correct answer was worth 1 point and the justification was worth an additional 2 points.

ii) 8MB 2-way set-associative cache with a Most Recently Used (MRU) replacement policy will always have better or equal cache performance when compared to a 4MB direct-mapped cache.

TRUE

FALSE

Why?

TRUE. A larger cache and higher degree of associativity cache will have a higher hit rate than a smaller direct cache (fewer capacity and conflict misses). The correct answer was worth 1 points and the justification was worth an additional 2 points.

Solutions NAME: _____

- iii) Suppose we have a multithreaded process that accesses (reads and writes) to an external database, and it avoids race conditions by having each thread acquire a lock global variable before accessing the database. Using this approach, we can run two instances of this process accessing the same database at the same time.

TRUE

FALSE

Why?

FALSE. A lock in one process instance will not affect the other process. The correct answer was worth 1 points and the justification was worth an additional 2 points.

- iv) If the Banker's algorithm will not approve a resource request, and the resource request is processed, then system necessarily will enter deadlock.

TRUE

FALSE

Why?

FALSE. Banker's algorithm only guarantees that a system is in a safe state. If a system is in an unsafe state, deadlock may or may not occur. The correct answer was worth 1 points and the justification was worth an additional 2 points.

Solutions NAME: _____

b. (15 points) Short answer.

- i) (4 points) Briefly, in two to three sentences, explain how having a combination of I/O-bound processes and CPU-bound processes maximizes system utilization.

We want to keep all parts of the system busy, so if you can have some processes using the CPU while others wait on I/O, then the system will be more utilized. In general, systems are underutilized (there are a lot of wasted cycles both on I/O devices as well as CPU).

- ii) (4 points) Briefly, in two to three sentences, explain how given a set of jobs and their arrival and execution times, you could quickly ($O(n)$ time) determine if FCFS would produce a schedule that is optimal in terms of average completion time.

If the arrival order and execution times would yield a FCFS schedule that is identical to SRTCF, then the FCFS schedule would be optimal.

- iii) (4 points) Briefly, in two to three sentences, explain when a hard real-time scheduling system, such as Earliest Deadline First, will **not** be able to meet all deadlines.

Solutions NAME: _____

If more tasks are submitted than the system has resources, then EDF will be unable to provide a schedule that meets deadlines.

- iv) (3 points) Early Intel processors such as the 8086 did not support dual-mode operation. Briefly, in two to three sentences, explain whether you can or cannot implement virtual memory and address spaces on such systems. If so, explain how. If not, explain why not.

No, special instructions are unprotected and any process could change address translation information. We also accepted a yes answer if you clearly explained how a virtual machine could be used to emulate the behavior of dual mode operation.

Solutions NAME: _____

2. (30 points total) Deadlock.

- a. (8 points) What are the four requirements for deadlock? Provide a brief one sentence explanation of each requirement.

#1

Mutual exclusion: at least one resource must be held in a non-sharable mode.

#2

Hold and wait: a process must be holding at least one resource and waiting to acquire additional resources that are currently being held by other processes.

#3

No preemption: a resource can be released only voluntarily by the process holding it.

#4

Circular wait: A set $\{P_0, P_1, \dots, P_n\}$ of waiting process must exist such that P_0 is waiting for a resource held by P_1 , P_1 is waiting for a resource held by P_2 , ... P_{n-1} is waiting for a resource held by P_n , and P_n is waiting for a resource held by P_0 .

Solutions NAME: _____

- b. (6 points) Consider the following resource allocation policy:
Requests for and releases of resources are allowed at any time. If a request for resources cannot be satisfied because the resources are not available, then we check any processes that are blocked, waiting for resources. If they have the desired resources, then these resources are taken away from them and are given to the requesting process. The vector of resources for which the waiting process is waiting is increased to include the resources that were taken away. If the resources needed by a blocked process become available, the process is put back on the ready queue.

For example, consider a system with three resource types and the vector *Available* initialized to (4,2,2).

- If process P0 asks for (2,2,1), it gets them.
 - If P1 asks for (1,0,1), it gets them.
 - Then, if P0 asks for (0,0,1), it is blocked (resource not available).
 - If P2 now asks for (2,0,0), it gets the available one (1,0,0) and one that was allocated to P0 (since P0 is blocked). P0's *Allocation* vector goes down to (1,2,1), and its *Need* vector goes up to (1,0,1).
- ii) With this resource allocation policy, can deadlock occur? If so, give an example. If not, which necessary condition cannot occur?
Deadlock cannot occur because preemption exists, which removes one of the preconditions for deadlock.

- iii) Can starvation occur?
Yes. A process may never acquire all the resources it needs if they are continuously preempted by a series of requests.

Solutions NAME: _____

- c. (7 points total) Consider a monitor implemented as follows:
 condition p, q; // p and q are initialized for you
 lock mutex = FREE;

```

void stop(void) {
    .....
    p.wait(&mutex);
    .....
    q.signal();
}

void go(void); {
    .....
    p.signal();
    q.wait(&mutex);
}
    
```

Suppose two processes P1 and P2 use the monitor in the following way:

```

P1: while (TRUE) {
    .....
    stop();
    .....
}

P2: while (TRUE) {
    .....
    go();
    .....
}
    
```

- i) (5 points) Show that deadlock can occur by listing an execution sequence ending in deadlock with explanation for each step. *Only one process can run at a time at each step – you do not have to fill in all the rows below.*

P1	P2	Description
	<i>go()</i>	<i>P2 calls go()</i>
	<i>p.signal()</i>	<i>P2 signals CV p. This signal is lost</i>
	<i>q.wait()</i>	<i>P2 blocks on CV q</i>
<i>stop()</i>		<i>P1 calls stop</i>
<i>p.wait()</i>		<i>P1 blocks on CV p</i>

Now, P1 is waiting for P2's signal (i.e., p.signal) and P2 is waiting for P1's signal (i.e., q.signal). Thus, we have a circular waiting, and, hence a deadlock..

- ii) (2 points) Can starvation occur?
Yes. Deadlock implies starvation.

Solutions NAME: _____

- d. (9 points) There are 5 lawyers sitting at a round dinner table. Lawyers repeat (forever) the following three things in order: (1) think, (2) eat, and (3) sleep. Each lawyer has a bowl of rice in front of them and a chopstick between each plate. **A lawyer must have two chopsticks to be able to eat!** A lawyer can only acquire the chopstick to the left of their plate and the chopstick to the right of their plate; they cannot reach across the table to use other chopsticks. When they are done eating, they put down (release) the two chopsticks for someone else to use. Below is code that attempts to solve this problem. Each lawyer will call this method with the parameter (0 to 4) indicating the lawyer's seat number. **Unfortunately, this code can deadlock.**
Your task is to fix this code so it cannot get stuck and so that more than one lawyer can eat at the same time – your solution cannot use any other synchronization primitives.

The constructor is called once in the main function, and then 5 lawyer threads are created, each one executing the Go() method. Think(), Eat(), and Sleep() are already written and you don't know how long they take to run.

```
Lawyers() {
    for (int i=0; i<=4; i++) chopstick[i] = new Semaphore(1);
}
void Go(int p) {
    while (1) {
        Think();
        // By having every other philosopher
        // get the chopsticks in the opposite
        // order, we break the cyclic dependency
        if (p%2 == 0) {
            // Grab left, then grab right

            chopstick[p].P();
            chopstick[(p + 1) % 5].P();

        } else {
            //---Grab right, then grab left
```

Solutions NAME: _____

```
        chopstick[(p + 1) % 5].P();
        chopstick[p].P();
    }
    Eat();
    chopstick[p].V();
    chopstick[(p + 1) % 5].V();
    Sleep();
    }
}
}
```

Note that we accepted any solution in which at least one lawyer acquired the chopsticks in reverse order also was correct.

Solutions NAME: _____

3. (25 points total) Memory Management.

a. (12 pts) Consider the following string of memory references in terms of requested page number: 1, 2, 3, 4, 5, 3, 4, 1, 6, 7, 8, 7, 8, 9, 7, 8, 9, 5, 4, 5, 4, 2.

Fill in the table below with using the FIFO, MIN and LRU (Least Recently Used) page replacement algorithms for memory with five frames by filling in the table with page numbers. The top row is the string of memory references, and each row contains the page numbers that reside in each physical frame F1 – F5 after each memory reference. For readability purposes, please only fill in the table entries that have changed and leave the unchanged entries blank. If several pages meet the replacement criteria, replace the one with the smallest frame number.

i) FIFO

	1	2	3	4	5	3	4	1	6	7	8	7	8	9	7	8	9	5	4	5	4	2
F1	1								6											5		
F2		2								7												2
F3			3								8											
F4				4										9								
F5					5														4			

What is the number of page faults, including compulsory faults? 12

Solutions NAME: _____

ii) MIN

	1	2	3	4	5	3	4	1	6	7	8	7	8	9	7	8	9	5	4	5	4	2
F 1	1								6	7												2
F 2		2												9								
F 3			3								8											
F 4				4																		
F 5					5																	

What is the number of page faults, including compulsory faults? 10

iii) LRU

	1	2	3	4	5	3	4	1	6	7	8	7	8	9	7	8	9	5	4	5	4	2
F 1	1																	5				
F 2		2							6										4			
F 3			3								8											
F 4				4										9								
F 5					5					7												2

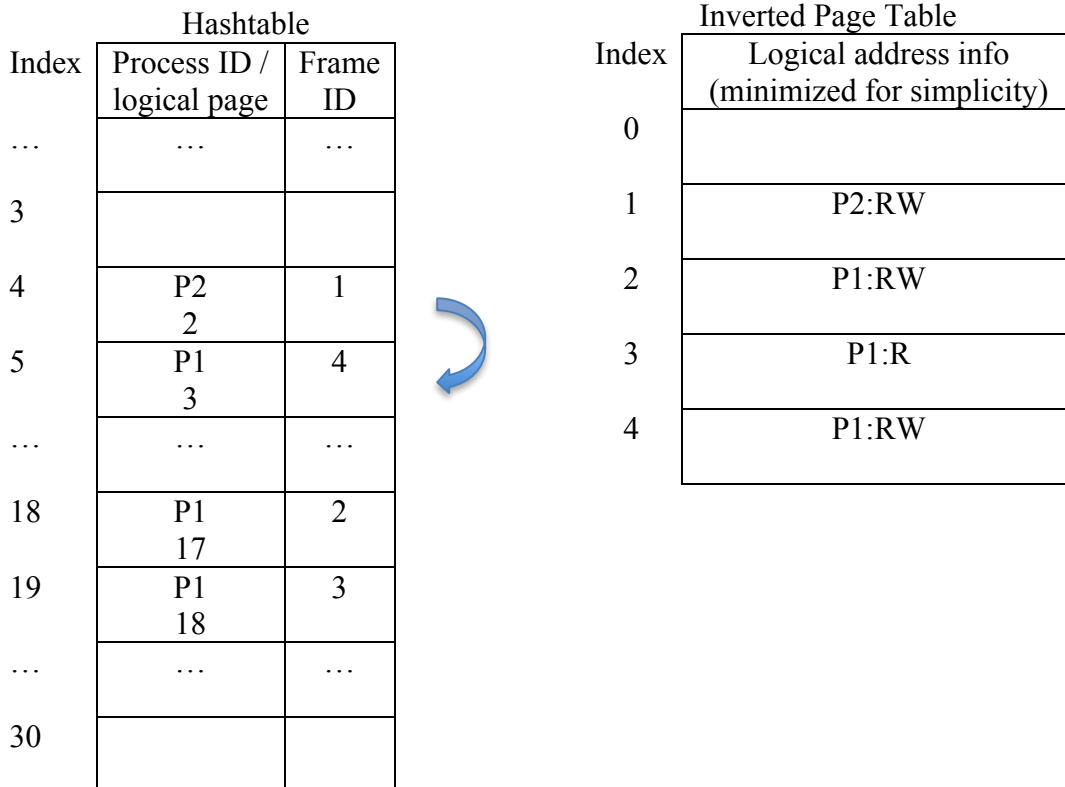
What is the number of page faults, including compulsory faults? 12

Solutions NAME: _____

- b. (4 points) Some operating systems provide a virtual copy operation: `VirtualCopy(proc, addr1, addr2, len)`, where data of length `len` starting at location `addr1` in the caller's address space is "copied" to process `proc`'s address space starting at location `addr2`. Assume that `addr1`, `addr2`, and `len` are multiples of the page size. Assuming a page-table-based virtual memory system, explain how we can efficiently implement `VirtualCopy` so as to minimize the amount of copying required. *Note that after `VirtualCopy` completes, the two processes will **not** share writes to the data. We could use copy-on-write: the page-table entries for the pages of the source (in the calling process starting at location `addr1` for `len` bytes) is copied into page table entries for process `proc`'s address space starting at location `addr2`. The pages are marked "copy on write" (if not done so already) and their reference counts are incremented by one (they would of course start at one when the pages are first allocated). The page-table entries in both processes are marked read-only. Now, while both processes merely read the pages, they continue to share them. But if either process attempts to modify the page, the attempt is trapped by the OS, a new page is allocated and the old is copied into it, the page table entry for the modifying process is set to point to the new page and is set to read-write with a reference count of 1, and the reference count of the original page is decremented by 1. If this reference count is now zero, its (single) page-table entry is changed to read-write.*

Solutions NAME: _____

- c. (9 points total) Inverted Page Tables. Assume we have a simple, demand paging environment, with no segmentation. Processes P 1 and P 2 both have logical memory addresses in the range 0 . . . 99, inclusive. The page size is 5. The hash table portion of the structure uses a hash function which simply calculates the index by adding the numerical portion of the process identifier and the logical page number. Note that this is a really bad hash function because it requires a large hash table, but for our exercise, it is sufficient. The hash table and the inverted page table shown are below.



Arrows in the hash table indicate chaining of entries when a hash conflict has occurred. Empty entries in the IPT indicate the associated physical frames are not allocated.

Solutions NAME: _____

- i) (6 points) Calculate the physical addresses for the following logical addresses where the number after the colon is the logical address (**not the logical page**). Assume the first page of a process is page number 0, and that a mechanism exists to find a free frame in memory. .

Show your work for these calculations for partial credit – what calculations and table lookups were necessary to determine the actual address in memory from the logical address.

(1) P1:17

Logical page $17/5 = 3$ offset = 2. hash table index = $3 + 1$, follow chain to entry 5. From lookup, logical frame = 4. Address = (frame)(frame size) + offset = $4 * 5 + 2 = 22$*

(2) P1:92

Logical page = 18. offset = 2. Hash table index = 19. From lookup logical frame = 3. Address = 17

(3) P1:111

Exception, memory reference out of range.

- ii) (3 points) When process P2 attempts to read logical address 97, what happens? Specifically, describe the changes in the data structures, and what the process perceives of these changes.

This causes a page fault. Frame 0 is allocated, and the entry in the IPT is updated with the logical address information for P2 frame 19. Next the

Solutions NAME: _____

hashtable entry at index 21 ($19 + 2$) is mapped from P2:19 to frame 0. After this processing is complete, the P2 resumes execution as before.

Solutions NAME: _____

4. (18 points total) Scheduling. Five processes A, B, C, D and E arrive in this order at the same time with the following CPU bursts.

Process	CPU Burst
A	7
B	2
C	3
D	6
E	4

Fill in the entries of the following table with turnaround times and total turnaround time for each indicated scheduling policy and each process. Ignore context switching overhead. Turnaround time is defined as the time a process takes to complete after it arrives.

Scheduling Policy	Turnaround Time					Total Turnaround Time
	A	B	C	D	E	
First Come First Served	7	9	12	18	22	68
Shortest Job First	22	2	5	15	9	53
Round-Robin (quantum=2)	22	4	13	21	17	77

Scratch space: