# **CS143A Sample Final**

# Part A:

# [True / False]

Mark the following (T-F) questions using T for True and F for False in the table below. Illegible answers will be considered incorrect. (Will have negative grading)

			T/F						
1.	Direct Memory Access (DMA) bypasses the cache or TLB to read directly from main memory.								
2.	In the pyramid-shaped memory hierarchy, registers appear above cache.								
3.	Multiple threads within the same process share heap memory but not stack memory.								
4.	In UNIX, the <i>exec</i> () system call creates a new child process, and the <i>fork</i> system call replaces the current process's program with a new program and executes it.								
5.	The round robin scheduling algorithm avoids the convoy effect.								
6.	Priority inversion is a solution to the priorit	<i>ty inheritance</i> problem.							
7.	The following pseudo-code for a process and TestAndSet function, without any other assumptions, solves the critical section problem between N processes:         function TestAndSet(boolean lock) {             boolean initial = lock;             lock = FALSE;             return initial;             }             Initial value of lock = TRUE;             lock = TRUE;             }        function Process(int N) {                  Repeat                  while(TestAndSet (lock)) no-op;                  critical section                  lock = TRUE;                  remainder section                  until FALSE                  }								
8.	The <i>Progress</i> requirement for the critical section problem states that: if no process is currently in the critical section and then one requests to enter the critical section, this requesting process will be granted access to the critical section in a finite period of time.								
9.	Paging breaks up a program's view of memory into logical blocks such as "main",         "global variables", and functions/procedures								
10.	<i>Late binding</i> of memory addresses enables the compiler to produce more efficient code than <i>early binding</i> .								
11.	Consider an OS that employs the <i>working set model</i> to track the number of recent page references. To avoid the problem of <i>thrashing,</i> it swaps out a process if the total # demand frames exceeds the # available frames.								



Part B:

### **Question B.1: Inter-process communication**

Briefly (1-2 sentences) explain the main difference between **shared memory** IPC and **message passing** IPC.

### **Question B.2: Round Robin and Time Quantum**

Consider N processes sharing the CPU in a round-robin fashion (N>=2). Assume that each context switch takes S ms and that each time quantum is Q ms. For simplicity, assume that processes never block on any event and simply switch between the CPU and the ready queue. Also, assume that a process is still in the ready queue while a context switch is happening.

a) What happens if Q is much smaller than S? What happens when  $Q \rightarrow \infty$ , i.e. is much larger than the maximum turnaround time of all the processes? Be brief (1-2 sentences max) in your answer.

b) If you use RR for scheduling, which of the three performance metrics (waiting, response, turnaround time) is more likely to be improved? Why (1-2 sentences max)?

# **Question B.3: CPU Scheduling Algorithms**

Consider the set of process:

Process ID	Arrival Time	Burst Time
P1	0	5
P2	0	10
P3	4	15
P4	18	10
P5	22	20

a) Draw the GANTT chart for the **Round Robin** (time quantum = 5) scheduling algorithm. Use the same implementation you used for the programming assignment i.e. the processes should always run in PID order. Assume there is no context-switch overhead. Show your work for partial credit.

b) Write your answer to the following performance metrics given your above GANTT charts. Show your work for partial credit.

Average Response time:

Average Waiting time:

Average Turnaround time:

# Part C:

# **Question C.1: Process Synchronization**

Recall the following process execution diagram:



Recall how we use the following semaphores to enforce the execution order above: s1=0; s2=0; s3=0;

P1: body; V(s1); V(s1); P2: P(s1); body; V(s2); P3: P(s1); body; V(s3);

P4: P(s2); P(s3); body;

Where the semaphores **s1**, **s2**, and **s3** are created with an initial value of 0.

(continued on next page)

Use the same semaphore notation shown above to describe how we can ensure the execution order of the following process execution graph:



Use **all of** the following semaphores in your answer:

### s1=0; s2=0; s3=0; s4=0; s5=0; s6=0;

(write your answer below here)

### **Question C.2: Readers / Writers**

Recall the Readers / Writers problem from lecture. Fill in the blanks below to properly complete the following pseudocode implementation:

#### Shared Data:

```
Semaphore wrt initialized to 1
Semaphore mutex initialized to 1
Integer read_count initialized to 0
```

#### Writer Process:

do { \_\_\_\_ . . . /\* writing is performed \*/ . . . } while (true); Reader Process: do { wait(mutex); read count++; if (read count == 1) signal(mutex); . . . /\* reading is performed \*/ . . . read count--; if (read count == 0)

} while (true);

### **Question C.3: Deadlocks & safe state**

Consider the following snapshot of a system:

Process		Alloc	ation		Мах				Available			
	A	В	С	D	A	В	С	D	A	В	С	D
P0	3	0	2	1	4	2	4	2	1	0	0	0
P1	0	1	0	1	0	2	2	2				
P2	1	2	0	0	3	2	1	0				
P3	0	1	1	2	1	1	1	2				
P4	0	0	1	1	1	0	2	1				

a) What is the content of the matrix *Need*?

Process	Need							
	А	В	С	D				
P0								
P1								
P2								
P3								
P4								

b) Is the system in a safe state? If yes, give a safe sequence of processes. If not, explain why the system is not in a safe state.

c) If a request from process P4 arrives for (1,0,0,0), can the request be granted immediately? Please state the reason.

# Part D:

### **Question D.1: Effective Access Time**

A computer keeps its page tables in memory. Memory access time is 100 nanoseconds (ns). Answer the following questions about the performance of this setup. Show your work.

a) What is the effective access time (i.e. reading a word in memory) with no caching and a two-level page table?

b) Consider the above scenario but with a TLB having a cache hit rate of 98%. If the TLB takes 20 ns to access, what is the effective access time of this setup when considering this TLB?

### **Question D.2: Bits needed for addressing**

Consider a system that uses a fixed-partition scheme, with equal partitions of size 2<sup>8</sup> bytes, and the main memory has 2<sup>16</sup> bytes. A process table is maintained with a pointer to the resident partition for each resident process. How many **bits** are required for the pointer in the process table? Show all your steps.

### **Question D.3: Segmentation**

Segment	Base (original)	Length	Base (after compaction)
0	100	300	
1	1400	600	
2	450	100	
3	3200	80	
4	2200	500	
5	3300	33	

Given the following segment table:

Given the **original** base addresses, what are the physical addresses for the following logical addresses? If it's an invalid address, just write "invalid". Note that (X, Y) => segment X, offset Y
 a) (0, 350)

b) (1, 599)

c) (2, 50)

d) (3, 81)

e) (4, 300)

f) (5, 0)

g) (5, 34)

2. If the above segments are the only segments in the system, i.e. there is only 1 process, and all segments are full, fill the right-most column of the base address table above after completing a disk compaction.

# **Question D.4: FIFO Page Replacement**

Consider the following page reference string:

e, c, b, e, a, g, d, c, e, g, d, a

Considering 4 frames, fill in the following table and then answer how many page faults would occur with the FIFO page replacement algorithm.

RS: reference string; F0: frame 0, F1: frame 1, etc.

Hint: all frames are initially empty, so your first unique pages will all cost one fault each.

Time	1	2	3	4	5	6	7	8	9	10	11	12
RS	е	С	b	е	а	g	d	С	е	g	d	а
F0												
F1												
F2												
F3												
Page fault?												

b) Total # page faults:

c) Briefly (1-2 sentences) explain *Belady's Anomaly* that can occur in FIFO Page Replacement.

# Part E:

# **Question E.1: Linked Allocation File System Operation**

Consider a file currently consisting of 10 blocks. Assume that the file control block and the new block information to be added are already in memory. Calculate how many disk I/O operations are required for the linked allocation strategy, if, for one block, the following conditions hold:

HINTS: 1) ignore disk I/O associated with the file control block. 2) each read and each write is an explicit disk I/O. 3) assume a pointer to the end of the list for linked allocation

- a. The block is added at the beginning.
- b. The block is added in the middle.
- c. The block is added at the end.
- d. The block is removed from the beginning.
- e. The block is removed from the middle.
- f. The block is removed from the end.

### **Question E.2: Indexed Allocation File System Operation**

Consider a *two-level indexed allocation* scheme that uses only a single top-level *index table*. If each block is *B* bytes and each index entry is *P* bytes, what is the maximum file length in bytes?