

University of California, Berkeley
College of Engineering
Computer Science Division – EECS

Fall 2013

Anthony D. Joseph and John Canny

Midterm Exam #1
October 21, 2013
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 | 8 | |
| 2 | 17 | |
| 3 | 15 | |
| 4 | 23 | |
| 5 | 18 | |
| 6 | 19 | |
| TOTAL | 100 | |

NAME: _____

1. (8 points total) True/False and Why? **CIRCLE YOUR ANSWER.**

i) A user-level process cannot modify its own page table entries.

TRUE

Why?

FALSE

ii) The scheduler is the part of an Operating System that determines the priority of each process.

TRUE

Why?

FALSE

iii) Shortest Remaining Time First is the best preemptive scheduling algorithm that can be implemented in an Operating System.

TRUE

Why?

FALSE

iv) The working set model is used to compute the *average* number of frames a job will need in order to run smoothly without causing thrashing.

TRUE

Why?

FALSE

NAME: _____

2. (17 points total) Deadlock.

- a. (11 points total) Recall the various deadlock detection and prevention algorithms we've discussed in this course, and consider the following snapshot of a system with five processes (P1, P2, P3, P4, P5) and four resources (R1, R2, R3, R4). There are no current outstanding queued unsatisfied requests.

Currently Available Resources

| R1 | R2 | R3 | R4 |
|----|----|----|----|
| 2 | 1 | 2 | 0 |

| Process | Current Allocation | | | | Max Need | | | | Still Needs | | | |
|---------|--------------------|----|----|----|----------|----|----|----|-------------|----|----|----|
| | R1 | R2 | R3 | R4 | R1 | R2 | R3 | R4 | R1 | R2 | R3 | R4 |
| P1 | 0 | 0 | 1 | 2 | 0 | 0 | 3 | 2 | 0 | 0 | 2 | 0 |
| P2 | 2 | 0 | 0 | 0 | 2 | 7 | 5 | 0 | 0 | 7 | 5 | 0 |
| P3 | 0 | 0 | 3 | 4 | 6 | 6 | 5 | 6 | 6 | 6 | 2 | 2 |
| P4 | 2 | 3 | 5 | 4 | 4 | 3 | 5 | 6 | 2 | 0 | 0 | 2 |
| P5 | 0 | 3 | 3 | 2 | 0 | 6 | 5 | 2 | 0 | 3 | 2 | 0 |

- i) (5 points) Is this system currently deadlocked, or can any process become deadlocked? Why or why not? If not deadlocked, give an execution order.

- ii) (3 points) If a request from a process P1 arrives for (0, 4, 2, 0), can the request be immediately granted? Why or why not? If yes, show an execution order.

NAME: _____

iii) (3 points) If a request from a process P2 arrives for (0, 1, 2, 0), can the request be immediately granted? Why or why not? If yes, show an execution order.

b. (6 points) Briefly *in at most three sentences each* describe two approaches to avoiding deadlock.

Approach #1:

Approach #2:

NAME: _____

3. (15 points total) Demand Paging

For each of the following page replacement policies, list the total number of page faults and fill in the contents of the page frames of memory after each memory reference.

a. (5 points) MIN page replacement policy:

| Reference | E | D | H | B | D | E | D | A | E | B | E | D | E | B | G |
|--------------------|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Page #1 | E | | | | | | | | | | | | | | |
| Page #2 | - | D | | | | | | | | | | | | | |
| Page #3 | - | - | H | | | | | | | | | | | | |
| Mark X for a fault | X | X | X | | | | | | | | | | | | |

Number of MIN page faults? _____

b. (5 points) LRU page replacement policy:

| Reference | E | D | H | B | D | E | D | A | E | B | E | D | E | B | G |
|--------------------|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Page #1 | E | | | | | | | | | | | | | | |
| Page #2 | - | D | | | | | | | | | | | | | |
| Page #3 | - | - | H | | | | | | | | | | | | |
| Mark X for a fault | X | X | X | | | | | | | | | | | | |

Number of LRU page faults? _____

c. (5 points) FIFO page replacement policy:

| Reference | E | D | H | B | D | E | D | A | E | B | E | D | E | B | G |
|--------------------|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Page #1 | E | | | | | | | | | | | | | | |
| Page #2 | - | D | | | | | | | | | | | | | |
| Page #3 | - | - | H | | | | | | | | | | | | |
| Mark X for a fault | X | X | X | | | | | | | | | | | | |

Number of FIFO page faults? _____

NAME: _____

4. (23 points) Memory management:

a. (5 points) Consider a memory system with a cache access time of 10ns and a memory access time of 110ns – assume the memory access time includes the time to check the cache. If the effective access time is 10% greater than the cache access time, what is the hit ratio H ? (**fractional answers are OK**)

b. (18 points total) Address Translation:

i) (5 points) Consider a machine with a physical memory of 8 GB, a page size of 8 KB, and a page table entry size of 4 bytes. How many levels of page tables would be required to map a 46-bit virtual address space if every page table fits into a single page? Be explicit in your explanation.

ii) (4 points) List the fields of a Page Table Entry (PTE) in your scheme.

NAME: _____

iii) (3 points) *Without a cache or TLB*, how many memory operations are required to read or write a single 32-bit word?

iv) (6 points) How much physical memory is needed for a process with three pages of virtual memory (for example, one code, one data, and one stack page)?

NAME: _____

5. (18 points total) Scheduling.

Assumptions: All timeslice-based algorithms have a timeslice of one unit; The currently running thread is not in the ready queue while it is running; An arriving thread is run at the beginning of its arrival time, if the scheduling policy allows it. Turnaround time is defined as the time a process takes to complete after it arrives.

Fill in ALL blanks in EACH table – each blank has an unambiguous answer.

For the missing schedulers, the possibilities are **SRTF, RR, and Priority**.

Priority is a preemptive scheduler.

Hint: Fill in the entry time (below) for Thread C first!

| Entry Times | |
|-------------|---|
| A | 1 |
| B | 2 |
| C | |
| D | 8 |

| Priorities | |
|------------|---|
| A | 3 |
| B | 4 |
| C | 5 |
| D | 6 |

| ↓ Current Time Scheduler → | Currently Scheduled Process | | |
|-------------------------------|-----------------------------|---|---|
| | FIFO | | |
| 1 | A | A | A |
| 2 | A | A | B |
| 3 | A | A | B |
| 4 | B | | |
| 5 | B | | |
| 6 | B | | |
| 7 | C | | |
| 8 | D | | |
| 9 | D | | |
| 10 | D | | |
| Avg Turnaround Time | 3.5 | | |

NAME: _____

6. (19 points) Concurrency Control.

You are the organizer of a gaming exhibit at the E3 Electronic Entertainment Expo. You want to allow the attendees to play your startup's new game demo. You model the attendees as threads, called *players*, and your job is to synchronize access to a single copy of the game, as follows:

- When a player arrives, he or she waits in a waiting area.
- Once there are 4 or more players waiting to play, you allow *exactly* 4 of them to leave the waiting area to begin playing. These four leave the waiting area and approach the game console.
- When a player reaches the console, the player waits until all four players are at the console, at which point all four players begin playing.
- Players may finish playing at any time. However, you cannot allow any new players to begin playing until all four players have left.
- You do not need to let players out of the waiting area in the order in which they arrived.
- You cannot assume that a player will ever finish playing.

You decide to solve this synchronization problem using two custom synchronization primitives, which have “barrier-like” semantics:

```
GameBarrier gb;
ConsoleBarrier cb;
```

Your task is to implement these synchronization primitives according to the specifications listed below. Each player thread has access to the two global barriers, and uses them in the following sequence:

```
void Player(ThreadID tid, GameBarrier gb, ConsoleBarrier cb) {
    gb.waitToPlay();
    cb.waitAtConsole();
    play();
    gb.donePlaying();
}
```

The `GameBarrier` can be in one of three states:

- `GAME_NOTREADY`: There are fewer than 4 players waiting to play. When the barrier is in this state, no player can progress beyond `waitToPlay()`.
- `GAME_FILLING`: There are (or were) at least 4 players waiting to play, and either we are in the first turn or else all four players from the prior turn have departed (via `donePlaying()`). When the barrier is in this state, a player can progress beyond `waitToPlay()`, and in fact four players must progress beyond this function. When exactly four players have progressed beyond this function, the barrier enters the `GAME_FILLED` state.
- `GAME_FILLED`: Four players have been sent to the console, and the turn is not over (meaning that the departure of all four from the console via `donePlaying()` has not yet taken place). When the barrier is in this state, no waiting player can progress beyond `waitToPlay()`.

NAME: _____

The `ConsoleBarrier` can be in one of two states:

- `CONSOLE_WAIT`: Four players have not yet arrived at the console in the current round. When the barrier is in this state, no player can progress beyond `waitAtConsole()`.
- `CONSOLE_ALLOW`: Four players have arrived in the current round. When the barrier is in this state, all four waiting players must progress beyond `waitAtConsole()`, after which the state reverts to `CONSOLE_WAIT`.

Your barrier will require the use of condition variables. Recall that condition variables provide three methods: `Condition.wait(Lock mutex)`, `Condition.signal()`, and `Condition.broadcast()`.

Locks provide two methods: `Lock.acquire()` and `Lock.release()`.

Note that part (but not all) of your work is to ensure that the barriers make the correct state transitions.

Below, where indicated, fill out the variables and methods for the `GameBarrier` and `ConsoleBarrier` objects.

```
public class GameBarrier {
    private static final int GAME_NOTREADY = 0;
    private static final int GAME_FILLING = 1;
    private static final int GAME_FILLED = 2;

    private Lock mutex;
    private int state;
    private Condition cv;

    /* SPECIFY ANY OTHER CLASS VARIABLES */

    public GameBarrier(Lock lock) {
        this.mutex = lock;
        this.state = GAME_NOTREADY;
        this.cv = new Condition();

        /* INITIALIZE ANY OTHER CLASS VARIABLES */
    }
}
```

NAME: _____

```
public void waitToPlay() {  
    /* YOU MUST FILL IN THIS FUNCTION */
```

```
}
```

```
public void donePlaying() {  
    /* YOU MUST FILL IN THIS FUNCTION */
```

```
}
```

NAME: _____

```
    }

    public class ConsoleBarrier {
        private static final int CONSOLE_WAIT = 0;
        private static final int CONSOLE_ALLOW = 1;

        private Lock mutex;
        private int state;
        private Condition cv;

        /* SPECIFY ANY OTHER CLASS VARIABLES */

        public ConsoleBarrier(Lock lock) {
            this.mutex = lock;
            this.state = CONSOLE_WAIT;
            this.cv = new Condition();

            /* INITIALIZE ANY OTHER CLASS VARIABLES */

        }

        public void waitAtConsole() {
            /* YOU MUST FILL IN THIS FUNCTION */

        }
    }
```