

Fill out the bottom of this page with your computing ID.
Write your computing ID at the top of each page in case pages get separated.

On my honor as a student I have neither given nor received aid on this exam.

1. (6 points) In the traditional Unix's monolithic kernel operating system design (like used by Linux and xv6), which of the following typically operations occur in kernel mode? **Select all that apply.**
- creating a new process when a process calls `fork`
 - checking whether a file exists when a process calls `open`
 - displaying text on the console when a process calls `printf`
 - assigning new memory to a process's address space when a process calls `malloc`
 - comparing strings when a process calls `strcmp`
 - switching from one process to another
2. (4 points) In a multi-level feedback queue scheduler, programs in the highest priority queue also _____ **Select all that apply.**
- tend to be less compute-intensive than programs at lower priorities
 - have longer timeslices than programs at lower priorities
 - will be demoted to a lower priority if they stop using the CPU before the end of their timeslice
 - will be demoted to a lower priority if they keep using the CPU up until the end of their timeslice
3. For **each** of the following statements, indicate whether they are true about
- a first-come first-served scheduler (FCFS)
 - a round-robin scheduler (RR)
 - a lottery scheduler, like the one required for the scheduler assignment (lottery)
 - a shortest job first scheduler without preemption (SJF)
 - a shortest remaining time first scheduler with preemption (SRTF)
 - a multi-level feedback queue scheduler (MLFQ)
 - Linux's Completely Fair Scheduler (CFS)
 - a strict priority scheduler with preemption (prio)

For each question, do not accounting for how users might adjust what threads they run (other than specifying scheduler parameters, like priorities) based on the scheduler in use.

For each question below, **select all that apply.**

- (a) (5 points) This kind of scheduler is subject to starvation, where some thread will **never** run, no matter how long it waits, when there are certain kinds of competing threads, even if the scheduler does not allow any one thread to run indefinitely.
- FCFS RR lottery SJF SRTF MLFQ CFS prio
- (b) (5 points) This kind of scheduler allows the relative importance of threads to be specified by the user or system administrator and used in the scheduling decision.
- FCFS RR lottery SJF SRTF MLFQ CFS prio
- (c) (5 points) This kind of scheduler is not actually possible to implement exactly on a typical operating system, because it requires knowing about the future behavior of programs.
- FCFS RR lottery SJF SRTF MLFQ CFS prio

4. (20 points) Consider the following incomplete C program that uses the POSIX API: (Assume all appropriate headers are included and ignore minor syntax errors.)

```

int main(void) {
    int fds[2] = {1, 1}; pid_t pid = (pid_t) -1; char c;

    _____; // LOCATION 1
    pid = fork();

    _____; // LOCATION 2
    if (pid == 0) { /* child: */

        _____; // LOCATION 3
        write(fds[1], "Hello!", 6);
    } else {

        _____; // LOCATION 4

        while (_____ > 0) { // LOCATION 5
            printf("%c", c);
        }
        waitpid(pid, NULL, NULL);
    }
    return 0;
}

```

To make this program output “Hello!”, one or more of the following statements need to be inserted into the blanks marked Location 1 through 5.

- A. pipe(fds)
- B. dup2(fds[1], STDOUT_FILENO)
- C. dup2(fds[0], STDIN_FILENO)
- D. close(fds[0])
- E. close(fds[1])
- F. close(STDIN_FILENO)
- G. read(fds[0], &c, 1)
- H. read(fds[1], &c, 1)
- I. write(fds[1], &c, 1)
- J. read(STDOUT_FILENO, &c, 1)

For each line above, write the letters A through J that would be appropriate to place in each location above. Some lines may be left blank and some letters may be left unused.

pipe(a) takes an array a of two file descriptors and creates a pipe where element 1 of the array is the write end of the pipe and element 0 is the read end.

dup2(from, to) takes two file descriptors from and to and assigns the file descriptor number to to reference the same open file that from does.

You may assume fork and waitpid and pipe will not fail, and that write never writes less than requested when passed an appropriate file descriptor.

5. (8 points) Consider the following C program that uses the pthreads API:

```
int x = 50;
void child(void *ignored_argument) {
    x = x + 10;
    printf("[child]_␣%d_␣", x);
    x = x + 10;
}
int main(void) {
    pthread_t child_thread;
    x = 100;
    pthread_create(&child_thread, NULL, child, NULL);
    x = x - 1;
    pthread_join(child_thread, NULL);
    printf("[main]_␣%d", x);
    return 0;
}
```

Assume that `pthread_create` and `pthread_join` do not fail, that reads and writes to `x` are atomic and are executed in each thread in the order specified in the program's code (not reordered by the compiler or processor), and that the program does not crash or fail to compile.

Which of the following outputs are possible? **Select all that apply.**

- [main] 110 [child] 119
- [child] 110 [main] 109
- [child] 109 [main] 120
- [child] 60 [main] 99
- [child] 110 [main] 119
- [child] 110 [main] 99
- [main] 110 [child] 109
- [child] 109 [main] 119

6. Consider implementing a “synchronized buffer” data structure. This structure has two operations:
- void Put(int value)
 - int Get()

Each time a thread calls Put(), it puts a value into a shared buffer and returns *after the item has been retrieved by another thread calling Get()*. Each item that is Put() into the buffer will be retrieved exactly once (assuming some thread eventually calls Get()).

If a thread calls Get() and an item is not yet available, it first waits until an item is available, then returns that item.

- (a) (30 points) Suppose we implement this data structure with mutexes and condition variables with the following incomplete code, where code to fill in is marked with a blank and a commented letter:

```
int current_value;
bool current_value_used = false; bool current_value_received = false;
pthread_mutex_t lock;
pthread_cond_t put_ready; pthread_cond_t put_done; pthread_cond_t get_ready;

void Put(int value) {
    pthread_mutex_lock(&lock);

    while (current_value_used) {
        pthread_cond_wait(&put_ready, &lock);
    }
    current_value = value;
    current_value_used = true;
    current_value_received = false;
    pthread_cond_signal(&get_ready);

    while ( /* A */ _____ ) {
        pthread_cond_wait(&put_done, &lock);
    }

    /* B */ _____;

    /* C */ _____;
    pthread_mutex_unlock(&lock);
}

int Get() {
    pthread_mutex_lock(&lock);

    while ( /* D */ _____ ) {
        pthread_cond_wait(&get_ready, &lock);
    }

    int result = current_value;
    current_value_received = true;

    /* E */ _____;
    pthread_mutex_unlock(&lock);
    return result;
}
```

Fill in each of the blanks (A, B, C, D, E) above. You may assume that there is code elsewhere that will correctly initialize the mutexes and condition variables.

- (b) Suppose we implement this synchronized buffer data structure with semaphores with the following code:

```
int current_value;
sem_t put_gate;
sem_t get_gate;

void Put(int value) {
    sem_wait(&put_gate); // wait = down or P operation
    current_value = value;
    sem_post(&get_gate); // post = up or V operation
}

int Get() {
    sem_wait(&get_gate);
    int result = current_value;
    sem_post(&put_gate);
    return result;
}
```

What should the initial values of the counting semaphores `put_gate` and `get_gate` be?

- i. (5 points) `put_gate`:

- ii. (5 points) `get_gate`:

7. (7 points) Suppose a multiprocessor system uses a cache coherency mechanism where:

- each processor has its own cache
- memory and these caches are all connected via a shared bus
- each block in the caches are Invalid or Modified (value is dirty and held by exactly one processor's cache) or Shared (value is not dirty and potentially held by another processor's cache)

Which of the following operations require a processor to send a message on the shared bus? **Select all that apply.**

- writing to a value which it currently has cached in the Modified state
- reading from a value which it currently does not have cached
- completing an atomic exchange operation on a value it currently has cached in the Modified state
- completing an atomic exchange operation on a value it currently has cached in the Shared state
- writing to a value which it currently does not have cached
- writing to a value which it currently has cached in the Shared state
- reading from a value which it currently has cached in the Shared state

8. Consider the following C++ code for a flight reservation system. The `makeReservation` function simultaneously tries to simultaneously reserve a seat (decrementing from a count of available seats) on several flights and returns true if it is successful; otherwise, if not enough seats were available, it returns false. Unfortunately, it has few bugs.

```
struct Flight { string name; int remainingSeats; pthread_mutex_t lock; };
bool makeReservation(vector<Flight*> flights) {
    for (int i = 0; i < flights.size(); ++i) {
        pthread_mutex_lock(&flights[i]->lock);
        if (flights[i]->remainingSeats == 0) {
            pthread_mutex_unlock(&flights[i]->lock);
            return false;
        }
    }
    for (int i = 0; i < flights.size(); ++i) {
        flights[i]->remainingSeats -= 1;
        pthread_mutex_unlock(&flights[i]->lock);
    }
    return true;
}
```

- (a) (10 points) When one of the flights does not have any seats left, the above `makeReservation` function can result in the program hanging later on — even if the program only has one thread. Explain **briefly** (at most one sentence) or give an example of how this could happen.

- (b) (8 points) Even when all flights have enough seats left, the above `makeReservation` function can result in the program hanging if it is called in a particular way from multiple thread. Describe two simultaneous calls that would cause this. You may use shorthand like `makeReservation([A, B, C])` to represent a call to `makeReservation` with a vector of flights *A*, *B*, and *C*.

- (c) (8 points) What is a way the hang described in the previous question (8(b)) be corrected? Explain briefly (at most one sentence).

9. Consider a system which, like x86, which has:

- 32-bit virtual addresses
- a 2-level page table
- 20-bit virtual page numbers
- 4096 (2^{12}) byte pages
- 4-byte page table entries
- 4096 byte page tables at each level

On this system, the first 10 bits of a virtual page number are used to look up the page table entry in the first level page table and the second 10 bits of are used to look up entries in the second level page table.

(a) (6 points) If the first-level page table entry, located at physical byte address `0x8010`, for virtual address `0x1020123` specifies that the second-level page table is located at physical byte address `0x4000`, then at what physical byte address is the second-level page table entry for this virtual address?

(b) (4 points) Suppose the second-level page table entry for virtual address `0x4444` is marked as not present (also known as ‘invalid’). Suppose the OS wants to make reads and writes to this address work in spite of this. What could the OS do from its fault handler when a page fault occurs from the program trying to write to address `0x4448` as part of a `pushq` (push 8 bytes on the stack) instruction? Suppose that the OS determines that writes to this address should go to physical memory address `0x8448`.

- update physical address `0x8448` location from the kernel, then return from the fault handler
- update the program’s saved stack pointer to contain `0x8450`, then return from the fault handler
- update the second-level page table entry to reference the page containing `0x8448`, then return from the fault handler
- update the second-level page table entry to reference the page containing `0x8448`, then restart the faulting program from the beginning of `main`