

INSTRUCTIONS

This is your exam. Complete it either at exam.cs61a.org or, if that doesn't work, by emailing course staff with your solutions before the exam deadline.

This exam is intended for the student with email address <EMAILADDRESS>. If this is not your email address, notify course staff immediately, as each exam is different. Do not distribute this exam PDF even after the exam ends, as some students may be taking the exam in a different time zone.

For questions with **circular bubbles**, you should select exactly *one* choice.

- You must choose either this option
- Or this one, but not both!

For questions with **square checkboxes**, you may select *multiple* choices.

- You could select this choice.
- You could select this one too!

You may start your exam now. Your exam is due at <DEADLINE> Pacific Time. Go to the next page to begin.

This is a **proctored, closed-book exam**. During the exam, you may **not** communicate with other people regarding the exam questions or answers in any capacity.

If there is something in a question that you believe is open to interpretation, please make a private post on Piazza asking for clarification. We will issue an announcement if we believe your question merits one.

We will overlook minor syntax errors in grading coding questions. You do not have to add necessary **#include** statements. For coding questions, the number of blank lines you see is a suggested guideline, but is not a strict minimum or maximum. There will be no limit on the length of your answer/solution.

a) Name

b) Student ID

c) Please read the following honor code: "I understand that this is a closed book exam. I hereby promise that the answers that I give on the following exam are exclusively my own. I understand that I am allowed to use one 8.5x11, double-sided, handwritten cheat-sheet of my own making, but otherwise promise not to consult other people, physical resources (e.g. textbooks), or internet sources in constructing my answers." Type your full name below to acknowledge that you've read and agreed to this statement.

1. (22 points) True/False

Please **EXPLAIN** your answer in **TWO SENTENCES OR LESS** (Answers longer than this may not get credit!). Also, answers without any explanation **GET NO CREDIT!**

a) (2 points)

1 Given any two threads A and B, thread A is always able to read data from thread B's stack.

- True
 False

2 Explain.

b) (2 points)

- 1 Given any two threads A and B, thread A is always able to read data from thread B's heap.
 True
 False

2 Explain.

c) (2 points)

- 1 A pointer will point to the same (virtual) memory address before and after a successful call to `fork()`.
 True
 False

2 Explain.

d) (2 points)

- 1 A system call is a function linked into a user-mode program that executes a “transition to kernel mode” instruction, then begins to execute arbitrary code with kernel privileges.
 True
 False

2 Explain.

e) (2 points)

1 A system call allows a process to execute certain instructions in the kernel's code in user mode.

- True
- False

2 Explain.

f) (2 points)

1 Without system calls, user programs cannot intentionally transfer control from a user program to the kernel.

- True
- False

2 Explain.

g) (2 points)

1 Without interrupts, the kernel cannot intentionally transfer control from a user program to the kernel.

- True
- False

2 Explain.

h) (2 points)

1 One advantage of matching every user-level thread with a kernel thread is that the resulting user-level thread can execute arbitrary instructions in kernel mode simply by using a message channel to communicate with the kernel thread.

- True
- False

2 Explain.

i) (2 points)

1 Every user-level thread must have a unique kernel-level thread in order to facilitate multithreading (and allow concurrency).

- True
- False

2 Explain.

j) (2 points)

1 After a thread is put on the queue for a condition variable, the thread will eventually get to run until completion.

- True
- False

2 Explain.

k) (2 points)

1 It is possible to create a critical section without disabling interrupts and without any locks, semaphores, or condition variables.

- True
- False

2 Explain.

l) (2 points)

1 Spinlocks—threads that spin in a loop until the lock is released—always perform worse than traditional blocking locks.

- True
- False

2 Explain.

m) (2 points)

1 To create a critical section in a uniprocessor, the user can just disable interrupts to avoid being preempted.

- True
- False

2 Explain.

n) (2 points)

1 Because `cond_wait()` releases the condition variable lock in its implementation, after the thread is woken up, the thread is explicitly responsible for calling `lock_acquire`.

- True
- False

2 Explain.

o) (2 points)

1 When using PintOS lists, a struct can only have one `list_elem`.

- True
- False

2 Explain.

p) (2 points)

1 In PintOS, every user program will eventually call the `exit` syscall (assuming that the kernel does not kill the program due to a fault).

- True
- False

2 Explain.

q) (2 points)

1 Since PintOS doesn't have a `fork` syscall, the `exec` syscall implicitly spawns a new child process.

True

False

2 Explain.

r) (2 points)

1 In PintOS, if `open` is called twice on the same file in the process (without anything else happening in between), the two file descriptors can be used interchangeably (similar to what would happen if `dup()` was used to generate the second file descriptor).

True

False

2 Explain.

s) (2 points)

1 In PintOS, the kernel stack can grow dynamically to handle deeply recursive computations.

True

False

2 Explain.

t) (2 points)

1 In PintOS, both user and kernel virtual memory is per process.

True

False

2 Explain.

u) (2 points)

1 In PintOS, when accessing an unmapped user virtual address, a user thread will page fault while a kernel thread will not.

True

False

2 Explain.

v) (2 points)

- 1 Say two processes on the same machine **open** the same file, and that there are no race conditions or errors in doing so – that is, **open** returns a valid file descriptor to both processes. If either process makes a call to **remove** this file from the filesystem, the other process cannot **write** to its file descriptor any longer.

True

False

- 2 Explain.

w) (2 points)

- 1 Assume `test.txt` doesn't already exist.

```
char buf[BUFSIZE];
memset(buf, 'a', 100);
int fd = open("test.txt", O_CREAT|O_RDWR);
lseek(fd, BUFSIZE, SEEK_SET);
lseek(fd, 0, SEEK_SET);
read(fd, buf, BUFSIZE);
```

Afterwards, `buf` will contain `BUFSIZE` a characters.

True

False

- 2 Explain.

x) (2 points)

- 1 After calling `exec`, a process can read from and write to the same file descriptors that the original process had open, even though the original process's address space has been replaced.

True

False

2 Explain.

y) (2 points)

1 Calling `pipe` on an array like `int fds[2]`; creates a read and write stream that local processes can use for unidirectional IPC. Under the hood, a pipe is implemented as a buffer in user space where the read end sits at the 0th index of the pipe and the write end sits at the 1st index of the pipe.

True

False

2 Explain.

z) (2 points)

1 When using sockets, the server process spins on a `while (true)` loop, awaiting new connections. When a connection arrives, the request is dequeued, the server calls `accept` and creates a new socketed connection, which returns a descriptor referencing the socket that the server process uses to service the client's future requests. Both the server and the client have the responsibility of closing their socket descriptors.

True

False

2 Explain.

aa) (2 points)

1 Each client hardware device (with a unique IP address), can have many individual sockets connected to the same server/port combination, such that the server can differentiate between them.

True

False

2 Explain.

2. (16 points) Multiple Choice

a) (3 pt) How many new processes could be created? Assume at least 2 individual calls to fork() succeed.

```
#include <sys/types.h>
```

```
#include <unistd.h>
```

```
int main() {  
    for (int i = 0; i < 3; i++) {  
        fork();  
    }  
    return 0;  
}
```

0

1

2

3

4

5

6

7

8

b) (3 pt) How many times could `fork()` be called? Assume at least 2 individual calls to `fork()` succeed.

```
#include <sys/types.h>
#include <unistd.h>

int main() {
    for (int i = 0; i < 3; i++) {
        fork();
    }
    return 0;
}
```

- 0
- 1
- 2
- 3
- 4
- 5
- 6
- 7
- 8

c) (2 pt) Which of the following need to be saved and restored on a context switch between two threads in the same process?

- Computational registers (integer and floating point)
- Page-table root pointer
- Buffered data in streams opened with `fopen`
- Contents of the file descriptor table
- Stack pointer register

d) (2 pt) Which of the following are valid ways for two threads in the same process to communicate?

- Reading/writing to a memory address in one thread's stack
- Reading/writing to a memory address in the process's heap
- Reading/writing to the same location in one thread's kernel's stack
- Reading/writing to an open `FILE*`
- Reading/writing to a pipe
- Reading/writing to sockets

- e) (2 pt) Which of the following are valid ways for two threads in different processes to communicate?
- Reading/writing to a memory address in one thread's stack
 - Reading/writing to a memory address in the process's heap
 - Reading/writing to the same location in one thread's kernel's stack
 - Reading/writing to an open FILE*
 - Reading/writing to a pipe
 - Reading/writing to sockets
- f) (2 pt) Which of the following are atomic operations?
- `x -= 100`
 - `y++`
 - `test&set`
 - `cond_wait`
 - `process_execute` (from Project 1!)
 - `pthread_mutex_init`
- g) (2 pt) Which of the following are true about condition variables?
- Allows for waiting in a critical section
 - Are commutative just like semaphores
 - Are sometimes subjected to spurious wakeup when dealing with Hoare semantics
 - Avoids busy waiting
- h) (2 pt) Which of the following are false about condition variables?
- Allows for waiting in a critical section
 - Are commutative just like semaphores
 - Are sometimes subjected to spurious wakeup when dealing with Hoare semantics
 - Avoids busy waiting
- i) (2 pt) When setting up the user stack in PintOS, which of the following properties must be satisfied? Select all that apply.
- The stack starts at the very top of the user virtual address space, in the page just below the virtual address `PHYS_BASE`.
 - The stack pointer, `esp`, must be aligned to a 16-byte boundary after pushing the return address.
 - The order of words (ex: `/bin/ls -l foo bar`) must be in the same order as they were inputted on the top of the stack.
 - The value of the return address doesn't have to be any specific value.

j) (2 pt) The kernel in PintOS must verify pointers before accessing user memory. Which of the following must be true about a valid pointer? Select all that apply.

- The pointer is located in mapped virtual memory.
- The pointer is located above `PHYS_BASE`.
- The pointer can be a null pointer.
- None of the Above

k) (2 pt) In Project 1 you had to implement the wait syscall. A common implementation was to create some shared struct with a reference count, or `ref_cnt`, that was discussed in section. What are some of the properties of this `ref_cnt`? Select all that apply.

- The `ref_cnt` can grow and shrink depending on the number of children a parent thread can have.
- The `ref_cnt` must be synchronized with some sort of lock as different threads can concurrently read and write to it.
- The `ref_cnt` could start at 2, to represent the starting relationship from the parent and from the child.
- None of the Above

l) (3 pt) Assume all calls to `read()`, `write()`, and `fork()` succeed. Suppose that `montymole.txt` was empty before running this block of code. The following code was run:

```
int main(int argc, char** argv) {
    int fd1 = open("montymole.txt", O_WRONLY);
    int fd2 = open("montymole.txt", O_WRONLY);
    write(fd1, "mole", 4);
    write(fd2, "whack", 5);
    write(fd2, "mole", 4);
    write(fd1, "mole", 4);
    write(fd1, "mole", 4);
    close(fd1);
    close(fd2);
}
```

Which of the following could be the content of “montymole.txt”?

- whacmolemole
- whacmolek
- molewhackmolemolemole
- whackmolemole
- Nothing

- m) (3 pt) Assume all calls to `read()`, `write()`, and `fork()` succeed. Suppose that `montymole.txt` was empty before running this block of code. The following code was run:

```
void new_thread(void* arg) {
    int* fd = (int*) arg;
    write(*fd, "whackwhackmole", 14);
}
int main(int argc, char** argv) {
    int fd1 = open("montymole.txt", O_WRONLY);
    pthread monty_mole;
    pthread_create(&monty_mole, NULL, new_thread, (void*) &fd1);
    write(fd1, "mole", 4);
    close(fd1);
}
```

Which of the following could be the content of “montymole.txt”?

- molewhackwhackmole
- whackwhackmolemole
- molewhackwhackwhack
- molewhack
- whackwhackmole

- n) (3 pt) Assume all calls to `read()`, `write()`, and `fork()` succeed. Suppose that `montymole.txt` was empty before running this block of code. The following code was run:

```
void fork_p(pid_t pid, int fd1) {
    if (pid == 0) {
        write(fd1, "whack", 5);
    }
    else {
        write(fd1, "mole", 4);
    }
}
int main(int argc, char** argv) {
    int fd1 = open("montymole.txt", O_WRONLY);
    fork_p(fork(), fd1);
    write(fd1, "mole", 4);
    close(fd1);
}
```

Which of the following could be the content of “montymole.txt”?

- molekmole
- whacmmole
- whackmole
- molemolee
- whackmolemole

o) (2 pt) Which of the following are true about files?

- We can write to any file (with write permissions) using `printf`.
- If I wanted to share a file descriptor with another thread, I would have to pass it as an argument because the new thread doesn't know the file descriptor of that file.
- If I wanted to share a `FILE*` with another thread, I would have to pass it as an argument because it doesn't have access to the `FILE*`.
- If I wanted to share a `FILE*` with a child process, I wouldn't have to do anything because the child has access to the `FILE*` as a result of the fork.
- None of the Above

p) (2 pt) Which of the following scenarios are valid (true and/or correct)?

- In Thread A, I make a new file and make a new thread, Thread B. I close the file in Thread B but will still be able to write to the new file in Thread A.
- If I make a new process (using `fork()`) to serve client requests in a web server, I can close both the socket and client connection in the child process.
- Two processes can communicate with each other using sockets.
- The `fdopen()` library function takes in a file descriptor and returns a file pointer to the same file. The resulting `FILE*` can have more permissions than requested with the original `open()` syscall for the file descriptor.
- None of the Above

3. (24 points) Short Answer

a) 1. Reader() {
 2. //First check self into system
 3. lock.acquire();
 4. while ((AW + WW) > 0) {
 5. WR++;
 6. okToRead.wait(&lock);
 7. WR--;
 8. }
 9. AR++;
 10. lock.release();
 11.
 12. // Perform actual read-only access
 13. AccessDatabase(ReadOnly);
 14.
 15. // Now, check out of system
 16. lock.acquire();
 17. AR--;
 18. if (AR == 0 && WW > 0)
 19. okToWrite.signal();
 20. lock.release();
 21. }
 22.
 23. Writer() {
 24. // First check self into system
 25. lock.acquire();
 26. while ((AW + AR) > 0) {
 27. WW++;
 28. okToWrite.wait(&lock);
 29. WW--;
 30. }
 31. AW++;
 32. lock.release();
 33.
 34. // Perform actual read/write access
 35. AccessDatabase(ReadWrite);
 36.
 37. // Now, check out of system
 38. lock.acquire();
 39. AW--;
 40. if (WW > 0){

```

41.     okToWrite.signal();
42. } else if (WR > 0) {
43.     okToRead.broadcast();
44. }
45. lock.release();
46. }

```

- 1 Above, we show the Readers-Writers example given in class. It used two condition variables, one for waiting readers and one for waiting writers. Suppose that all of the following requests arrive in very short order (while R1 and R2 are still executing):

Incoming stream: R1,R2,W3,R4,W5,W6,R7,R8,R9,W10,R11,W12,W13,R14,R15,W16,R17

In what order would the above code process the above requests? If you have a group of requests that are equivalent (unordered), indicate this clearly by surrounding them with braces '{}'. You can assume that the wait queues for condition variables are FIFO in nature (i.e. signal() wakes up the oldest thread on the queue).

- 2 Above, we show the Readers-Writers example given in class. It used two condition variables, one for waiting readers and one for waiting writers. Suppose that all of the following requests arrive in very short order (while R1 and R2 are still executing):

Incoming stream: R1,R2,R3,W4,W5,R6,W7,R8,W9,R10,W11,W12,R13,R14,R15,W16,R17

In what order would the above code process the above requests? If you have a group of requests that are equivalent (unordered), indicate this clearly by surrounding them with braces '{}'. You can assume that the wait queues for condition variables are FIFO in nature (i.e. signal() wakes up the oldest thread on the queue).

- 3 Above, we show the Readers-Writers example given in class. It used two condition variables, one for waiting readers and one for waiting writers. Suppose that all of the following requests arrive in very short order (while R1 and R2 are still executing):

Incoming stream: R1,R2,R3,R4,W5,W6,W7,W8,R9,R10,R11,W12,R13,W14,R15,W16,R17

In what order would the above code process the above requests? If you have a group of requests that are equivalent (unordered), indicate this clearly by surrounding them with braces '{}'. You can assume that the wait queues for condition variables are FIFO in nature (i.e. signal() wakes up the oldest thread on the queue).

- 4 Above, we show the Readers-Writers example given in class. It used two condition variables, one for waiting readers and one for waiting writers. Suppose that all of the following requests arrive in very short order (while R1 and R2 are still executing):

Incoming stream: R1,R2,R3,R4,R5,W6,W7,W8,W9,R10,R11,W12,W13,R14,W15,R16,R17

In what order would the above code process the above requests? If you have a group of requests that are equivalent (unordered), indicate this clearly by surrounding them with braces '{}'. You can assume that the wait queues for condition variables are FIFO in nature (i.e. signal() wakes up the oldest thread on the queue).

- 5 Above, we show the Readers-Writers example given in class. It used two condition variables, one for waiting readers and one for waiting writers. Suppose that all of the following requests arrive in very short order (while R1 and R2 are still executing):

Incoming stream: R1,R2,R3,W4,W5,W6,W7,R8,R9,W10,R11,W12,R13,W14,R15,R16,R17

In what order would the above code process the above requests? If you have a group of requests that are equivalent (unordered), indicate this clearly by surrounding them with braces '{}'. You can assume that the wait queues for condition variables are FIFO in nature (i.e. signal() wakes up the oldest thread on the queue).

- 6 Above, we show the Readers-Writers example given in class. It used two condition variables, one for waiting readers and one for waiting writers. Suppose that all of the following requests arrive in very short order (while R1 and R2 are still executing):

Incoming stream: R1,R2,W3,W4,R5,W6,W7,W8,R9,W10,R11,R12,R13,R14,R15,R16,W17

In what order would the above code process the above requests? If you have a group of requests that are equivalent (unordered), indicate this clearly by surrounding them with braces '{}'. You can assume that the wait queues for condition variables are FIFO in nature (i.e. signal() wakes up the oldest thread on the queue).

- 7 Above, we show the Readers-Writers example given in class. It used two condition variables, one for waiting readers and one for waiting writers. Suppose that all of the following requests arrive in very short order (while R1 and R2 are still executing):

Incoming stream: R1,R2,R3,R4,R5,W6,W7,W8,R9,W10,W11,W12,R13,R14,R15,W16,R17

In what order would the above code process the above requests? If you have a group of requests that are equivalent (unordered), indicate this clearly by surrounding them with braces '{}'. You can assume that the wait queues for condition variables are FIFO in nature (i.e. signal() wakes up the oldest thread on the queue).

8 Explain how you got your answer.

A large, empty rectangular box with a thin black border, intended for the student to write their explanation for question 8.

9 (3 pt) Let us define the logical arrival order by the order in which threads first acquire the monitor lock. Suppose that we wanted the results of reads and writes to the database to be the same as if they were processed one at a time – in their logical arrival order; for example:

Incoming stream: W1,W2,R3,R4,R5,W6,W7,R8,R9,R10

Would be processed as: W1,W2,{R3,R4,R5},W6,W7,{R8,R9,R10}

regardless of the speed with which these requests arrive. Explain why the above algorithm does not satisfy this constraint.

The following program is run:

```
b) #include <unistd.h>
#include <pthread.h>
#include <stdio.h>
#include <sys/wait.h>

void *helper(void *arg) {
    int *number = (int *) arg;
    (*number)++;
    return NULL;
}

int main(int argc, char *argv[]) {
    int data = 0;
    pthread_t thread;
    pthread_create(&thread, NULL, helper, (void *) &data);
    pthread_join(thread, NULL);
    printf("data is %d\n", data);
    data = 0;
    pid_t cpid = fork();
    if (cpid == 0) {
        helper((void *) &data);
        return 0;
    } else {
        wait(NULL);
    }
    printf("data is %d\n", data);
}
```

1 (2 pt) What does this program output? (There should be 2 lines of output.)

2 (2 pt) Explain why the two lines of output are either different or the same.

c) (2 pt) Assume that you are using a monitor for synchronizing some shared data. Explain why you might want to wrap a while loop around a `cond_wait()` statement even in a system with **Hoare** scheduling.

d) (2 pt) Explain why **Hoare** scheduling for a monitor might result in less efficient use of the processor cache than **Mesa** scheduling.

e) (2 pt) Assume that you are using a monitor for synchronizing some shared data. Explain why you need a while loop around a check for a condition that puts a thread to sleep on a condition variable when using **Mesa** scheduling.

- f) (2 pt) Assume that you are using a monitor for synchronizing some shared data. Explain why you usually do not need a while loop around a check for a condition that puts a thread to sleep on a condition variable when using **Hoare** scheduling.

- g) (2 pt) Explain in 1-2 *short* sentences when we would **NOT** need to incorporate reference counts in structs that are shared between processes in PintOS.

- h) (2 pt) John (M) and Frank decide the implementation of threads and processes in PintOS is too trivial. They are now interested in adding the ability to create multithreaded processes. Recommend to them whether they should base their implementation on calling functions from the `pthread` library (i.e. `pthread_create`) or take an alternative approach. Defend your recommendation in no more than 3 sentences.

- i) (2 pt) In 3-4 sentences, explain how you would implement a `fork()` syscall in PintOS. Specifically, discuss how you would implement two different return values from the same system call.

j) (2 pt) In section 1, we reviewed how the syscall handler protects the kernel from corrupt or malicious user code. Here was the solution:

1. The syscall number is used to index into a vector mapping number of fixed syscall handlers -- this prevents the user from getting the kernel to run user code in kernel mode.
2. The handler copies the user arguments from the user registers or stack onto the kernel stack -- this protects the kernel from malicious code on the user stack from evading checks.
3. All of the arguments are validated before the syscall is executed -- this protects the kernel from errors in the user arguments like invalid or null pointers.
4. Results from the syscall are copied back into user memory so the caller has access to them.

Discuss why steps 2 and 3 can or cannot be switched in PintOS in no more than 3 sentences.

k) (2 pt) Describe two non-trivial differences between a function that reads in an entire file into a buffer using `fread` vs. `read`.

l) (2 pt) Neil is working on HW2 and is implementing the pipes portion of the homework. As part of his implementation, he has decided that the standard output of a parent process will be the standard input to its child process. Neil first calls `fork()`, then the child process calls `pipe()` to facilitate IPC with the parent process. Is this a good idea or not? Why?

- m) (2 pt) Suppose Taj was coding a multithreaded echo server (similar to the one in Section 4). He is making a new thread to serve a new client request, and thinks it would be a good idea to pass in the `server_socket` and the `client_socket` to the new thread. Then, once the thread is done resolving the request, the thread would close the `server_socket` as well as the `client_socket`. Is this a good idea? Explain your reasoning.

- n) (2 pt) When would it be better to use threads over processes to serve client requests?

- o) (2 pt) When would it be better to use processes over threads to serve client requests?

- p) (2 pt) How does a web server keep communication with various simultaneous clients separate (so that responses do not get interleaved)?

4. (22 points) Synchronization Primitives (Or Not?)

Recall from lecture that disabling interrupts is a valid way to gain mutual exclusion to the processor and hence all data in RAM on a uniprocessor system. However, disabling interrupts around critical sections comes with some downsides.

- a) (2 pt) Explain two downsides of disabling interrupts around critical sections on a uniprocessor.

- b) To solve this problem, we use disabling interrupts sparingly and leverage the ability to put threads to sleep while they wait for access to the critical section. Given functions `intr_disable()` and `intr_enable()`, which disable interrupts and enable interrupts, respectively, implement a semaphore that (a) does not busy wait on `sema_down` operations and (b) does not keep interrupts disabled for a large number of operations. You are also allowed to use a PintOS list in your solution as well as the following functions:

- `thread_current()`: gives a reference to the current thread. Assume this thread has an “elem” list element that can be put onto PintOS lists
- `thread_block()`: blocks the current thread and chooses a new thread to run
- `thread_block_and_intr_enable()`: blocks the current thread, chooses a new thread to run, and then enables interrupts
- `thread_unblock(thread)`: unblocks the given thread and puts it on the ready-queue

The solution for these questions requires no more than 6 lines each. You may find the PintOS function signatures at the end of the exam useful.

- 1 (1 pt)

```
typedef struct semaphore {
    -----
    -----
} sema_t;
```

2 (1 pt)

```
void sema_init(sema_t *sema, int value) {  
    -----  
    -----  
}
```

3 (2 pt)

```
void sema_down(sema_t *sema) {  
    -----  
    -----  
    -----  
    -----  
    -----  
    -----  
    -----  
}
```

4 (2 pt)

```
void sema_up(sema_t *sema) {  
    -----  
    -----  
    -----  
    -----  
    -----  
}
```



Semaphores are a powerful synchronization primitive. They can be used to implement both locks and condition variables. See the implementation of locks below:

```
c) typedef struct lock {  
    sema_t sema;  
} lock_t;  
  
void lock_init(lock_t *lock) {  
    sema_init(&lock->sema, 1);  
}  
  
void lock_acquire(lock_t *lock) {  
    sema_down(&lock->sema);  
}  
  
void lock_release(lock_t *lock) {  
    sema_up(&lock->sema);  
}
```

Semaphores can also be used to implement condition variables (and therefore monitors)! Implement condition variables via semaphores below. For simplicity, we allow you to assume that semaphores can be added to PintOS lists.

1 (1 pt)

```
typedef struct condition_variable {  
    -----  
} cv_t;
```

2 (1 pt)

```
void cond_init(cv_t *cv) {  
    -----  
}
```

3 (2 pt)

```
void cond_wait(cv_t *cv, lock_t *lock) {  
    -----  
    -----  
    -----  
    -----  
    -----  
    -----  
}
```

4 (2 pt)

```
void cond_signal(cv_t *cv, lock_t *lock) {  
    -----  
    -----  
}
```



5 (2 pt)

```
void cond_broadcast(cv_t *cv, lock_t *lock) {
    -----
    -----
}
```

We have now implemented all the synchronization primitives from enabling and disabling interrupts, where interrupts are not disabled for too long and the processor does not busy wait! You did it! You now have the ability to write synchronized data structures that are thread-safe!

Let us now forget everything we just did, and implement a data structure without synchronization primitives. We will implement a thread-safe queue via a circular array buffer without any synchronization primitives at all. For simplicity, assume that the following operations never occur:

- d)
 - Enqueueing when the queue is already full
 - Dequeueing when the queue is empty

Consider the single-threaded implementation of the queue via a circular array buffer:

```
# define MAX_SIZE 128

typedef struct circular_array_buffer {
    int FRONT;
    int BACK;
    int values[MAX_SIZE];
} queue_t; // U r a queue_t \pi

void queue_init(queue_t *queue) {
    queue->FRONT = queue->BACK = 0;
    for (int i = 0; i < MAX_SIZE; i++)
        queue->values = 0;
}

void queue_push(queue_t *queue, int value) {
    queue->values[queue->FRONT++ % MAX_SIZE] = value;
}

int queue_pop(queue_t *queue) {
    return queue->values[queue->BACK++ % MAX_SIZE]
}
```


If multiple threads access this data structure concurrently, there will (clearly) be race conditions. We can solve these race conditions by locking accesses to the queue and granting mutual exclusion. In this case, in the presence of contention, one or more threads would sleep while another thread is in the critical section. Instead, we implement a non-blocking, thread-safe version of the queue without synchronization primitives.

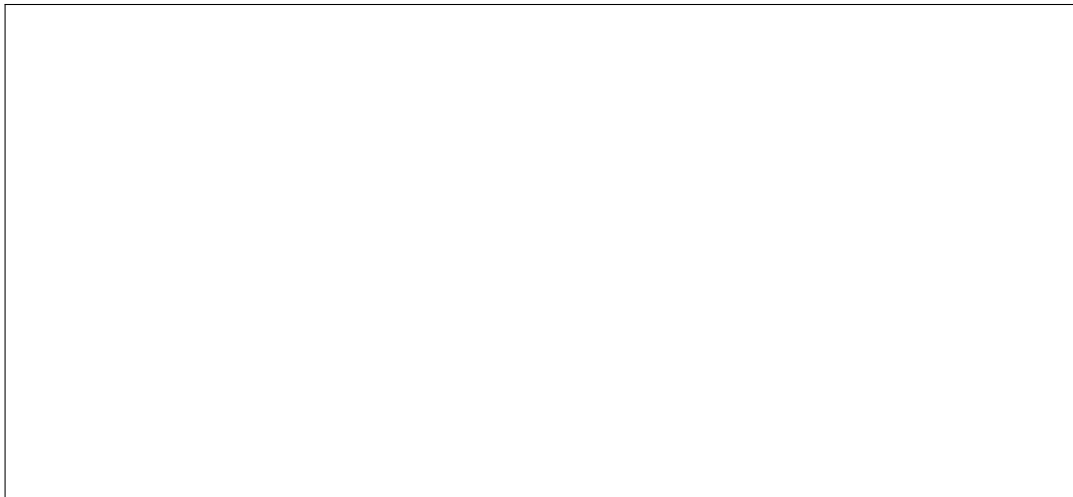
Armed with only the `compare&swap` and `compare&swap2` atomic operations (defined below) implement the `queue_push` and `queue_pop` operations to be thread-safe.

```
// Typical compare&swap operation
compare&swap(&addr, reg1, reg2)
    if (M[addr] == reg1)
        M[addr] = reg2
        return success
    else
        return failure

// Special compare&swap operation that allows two memory updates
compare&swap2(&addr1, reg1, reg2, &addr2, reg3)
    if (M[addr1] == reg1)
        M[addr1] = reg2
        M[addr2] = reg3
        return success
    else
        return failure
```

1 (3 pt)

```
void queue_push(queue_t *queue, int value) {
    -----
    -----
    -----
    -----
    -----
}
```



2 (3 pt)

```
int queue_pop(queue_t *queue) {  
    -----  
    -----  
    -----  
    -----  
    -----  
}
```

**5. (16 points) PintOS Pipes**

In this question you will be implementing support for a highly-simplified version of pipes within PintOS. Specifically, your pipes implementation will be simplified in the following ways:

- Each process has **exactly** one pipe.
- A process may only write to a pipe belonging to one of its direct children, using the `write_pipe` syscall.
- A process may only read from its own pipe, using the `read_pipe` syscall.
- Attempting to read from an empty pipe will result in the process being blocked until data is available (this is also true of normal Linux pipes). However, writing to a full pipe **should not** block - for the sake of keeping this question short, we are assuming a process will never write more data to a pipe than the pipe has capacity for.
- Pipes are not given file descriptors (hence the dedicated syscalls), and as such none of the conditions surrounding what happens if a write/read descriptor for a pipe is still open apply to this problem.

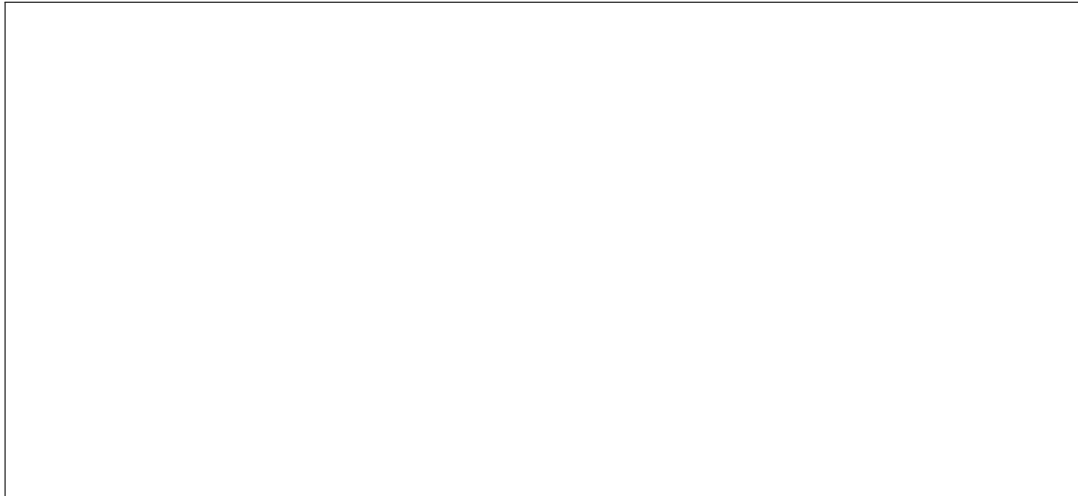
a) Part 1: Maintaining pipe state

Fill in the blanks below with whatever additional state you need to implement support for pipes. (Hint: some of this should be quite similar to what you did in project 1 in order to implement the `wait` syscall)

1 (2 pt) Inside threads/thread.h:

```
#define PIPE_BUFFER_SIZE 256

/*
 * Assume the pipe struct is properly reference counted
 * and therefore properly freed, in a manner similar to project 1
 * and the example shown with "Shared Data" in Section 2.
 *
 * Also assume that adding a child's pipe to its parent's
 * child_pipes list is taken care of for you.
 */
struct pipe {
    /* Reference count handling fields (not shown) */
    tid_t tid;
    uint8_t buffer[PIPE_BUFFER_SIZE];
    struct list_elem elem; /* List element for child_pipes */
    -----
    -----
    -----
    -----
};
```



2 (1 pt) Inside threads/thread.h:

```
struct thread {  
    tid_t tid;  
    struct list child_pipes;  
    /* Fields irrelevant to this problem (not shown) */  
    -----  
};
```



3 (3 pt) Inside `userprog/process.c`:

Fill in the below blanks with code that initializes the state that you added above.

```
static void start_process(/* Omitted */) {
    struct thread* cur = thread_current();
    /* Initialize interrupt frame, load executable, initialize
     * other state in thread struct (not shown).
     */
    -----
    -----
    -----
    -----
}
```

Part 2: Implementing the `read_pipe` and `write_pipe` syscalls

Assume that below functions are called when the `read_pipe` and `write_pipe` syscalls are made, respectively. **The below functions should work even if the `read_pipe` and `write_pipe` syscalls are made multiple times during a thread's lifetime.**

Inside `userprog/syscall.c`:

```
b) /*
   * Kills the current thread if pointer is not a valid pointer
   * within the user address space
   */
static void verify(void* pointer) {
    /* Code not shown */
}

/*
 * Returns the pipe belonging to child_tid if child_tid the tid
 * of a child of the current process, and returns NULL otherwise
 */
static struct pipe* helper(tid_t child_tid) {
    /* Code not shown */
}
```




2 (4 pt)

```
/*
 * Writes `length` bytes from `buffer` to thread `child_tid`'s pipe
 * Assume that there is always enough space in the pipe's buffer
 * to write `length` bytes without overflowing it.
 */
static int sys_write_pipe(tid_t child_tid, uint8_t* buffer, unsigned length) {
    verify(buffer);
    verify(buffer + length - 1);
    struct thread* cur = thread_current();
    -----
    -----
    -----
    -----
    -----
    -----
    -----
    -----
    return 0;
}
```

6. Reference Sheet

/***** Threads *****/


```

int pthread_create(pthread_t *thread, const pthread_attr_t *attr,
void *(*start_routine) (void *), void *arg);
int pthread_join(pthread_t thread, void **retval);
int pthread_mutex_init(pthread_mutex_t *restrict mutex,
const pthread_mutexattr_t *restrict attr);
int pthread_mutex_lock(pthread_mutex_t *mutex);
int pthread_mutex_unlock(pthread_mutex_t *mutex);
int sem_init(sem_t *sem, int pshared, unsigned int value);
int sem_post(sem_t *sem);
int sem_wait(sem_t *sem);

/***** Processes *****/
pid_t fork(void);
pid_t wait(int *status);
pid_t waitpid(pid_t pid, int *status, int options);
int execv(const char *path, char *const argv[]);

/***** High-Level I/O *****/
FILE *fopen(const char *path, const char *mode);
FILE *fdopen(int fd, const char *mode);
size_t fread(void *ptr, size_t size, size_t nmemb, FILE *stream);
size_t fwrite(const void *ptr, size_t size, size_t nmemb, FILE *stream);
int fclose(FILE *stream);

/***** Sockets *****/
int socket(int domain, int type, int protocol);
int bind(int sockfd, struct sockaddr *addr, socklen_t addrlen);
int listen(int sockfd, int backlog);
int accept(int sockfd, structure sockaddr *addr, socklen_t *addrlen);
int connect(int sockfd, struct sockaddr *addr, socklen_t addrlen);
ssize_t send(int sockfd, const void *buf, size_t len, int flags);

/***** Low-Level I/O *****/
int open(const char *pathname, int flags);
ssize_t read(int fd, void *buf, size_t count);
ssize_t write(int fd, const void *buf, size_t count);
int dup(int oldfd);
int dup2(int oldfd, int newfd);
int pipe(int pipefd[2]);
int close(int fd);

/***** PintOS *****/
void list_init(struct list *list);
struct list_elem *list_head(struct list *list)
struct list_elem *list_tail(struct list *list)
struct list_elem *list_begin(struct list *list);
struct list_elem *list_next(struct list_elem *elem);
struct list_elem *list_end(struct list *list);
struct list_elem *list_remove(struct list_elem *elem);
bool list_empty(struct list *list);
#define list_entry(LIST_ELEM, STRUCT, MEMBER) ...
void list_insert(struct list_elem *before, struct list_elem *elem);
void list_push_front(struct list *list, struct list_elem *elem);
void list_push_back(struct list *list, struct list_elem *elem);
void sema_init(struct semaphore *sema, unsigned value);

```

```
void sema_down(struct semaphore *sema);
void sema_up(struct semaphore *sema);
void lock_init(struct lock *lock);
void lock_acquire(struct lock *lock);
void lock_release(struct lock *lock);
void *memcpy(void *dest, const void *src, size_t n);
void *memmove(void *dest, const void *src, size_t n);
```