CS 162 Summer 2022 — Operating Systems and Systems Programming

Midterm Exam

INSTRUCTIONS

Please do not open this exam until instructed to do so.

Do not discuss exam questions for at least 24 hours after the exam ends, as some students may be taking the exam at a different time.

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!

Preliminaries

This is a proctored, closed-book exam. You are allowed 1 page of notes (both sides). You may not use a calculator. You have 110 minutes to complete as much of the exam as possible. Make sure to read all the questions first, as some are substantially more time-consuming.

If there is something about the questions you believe is open to interpretation, please ask us about it.

We will overlook minor syntax errors when grading coding questions. **There is a reference sheet at the end of the exam that you may find helpful.**

(a) Name

(b) Student ID

(c) Please read the following honor code: "I understand this is a closed-book exam. I promise the answers I give on this exam are 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 the internet in constructing my answers."

Write your full name below to acknowledge that you've read and agreed to this statement.

1. (12 points) True/False

Please explain your answer in **TWO SENTENCES OR LESS**. Longer explanations may get no credit. Answers without an explanation **GET NO CREDIT**.

(a) (2 points)

i. In Pintos the virtual address `0xc0000008` refers to the same physical address no matter which process we are in.

◯ True.

◯ False.

ii. Explain.

(b) (2 points)

i. After a multi-threaded process returns from `main`, the process will continue running until all threads call `pthread_exit` (assuming no errors).

◯ True.

◯ False.

ii. Explain.

(c) (2 points)

i. A thread can modify the stack variables of another thread in the same process.

◯ True.

◯ False.

ii. Explain.

(d) (2 points)

i. It is possible to create a lock without trapping into the kernel.

◯ True.

◯ False.

ii. Explain.

(e) (2 points)

i. `fork()` is slow in Linux because it has to copy the entire address space of a process.

◯ True.

◯ False.

ii. Explain.

(f) (2 points)

i. To maximize throughput, one should use FCFS instead of SRTF.

◯ True.

◯ False.

ii. Explain.

2. (18 points) Multiple Choice

Choose ALL that apply.

(a) (2 pt) Which of the following x86 instructions can you **NOT** execute in user mode?

- ☐ `pushl`
- ☐ `addl $16, %esp`
- ☐ `iret`
- ☐ `int $0x30`
- ☐ None of the above.

(b) (2 pt) Which of the following are part of the kernel?

- ☐ C standard library
- ☐ Device drivers
- ☐ Scheduler
- ☐ Shell
- ☐ None of the above.

(c) (2 pt) How does the Pintos kernel find the `struct thread` of the currently running kernel thread?

- ☐ The address is stored in an `intr_frame`.
- ☐ The kernel scans a global array of `struct thread`s.
- ☐ The address is stored in the upper 12 bytes of `%eax`.
- ☐ The kernel rounds down `%esp` to the nearest page.
- ☐ None of the above.

(d) (2 pt) Which of the following will ALWAYS lead to a transfer from user mode to kernel mode?

- ☐ A process opens a file.
- ☐ A process adds two large numbers, leading to integer overflow.
- ☐ A process tries to dereference a `NULL` pointer.
- ☐ A process reads a file with `fread`.
- ☐ None of the above.

(e) (2 pt) Which of the following are TRUE about syscalls?

- ☐ Forked processes share a single address space.
- ☐ Pipes are two-way communication channels between processes on the same physical machine.
- ☐ Open file descriptors are preserved across the (Linux) `exec` syscall.
- ☐ The `dup2` syscall allows a file descriptor to point to multiple entries in the kernel file table.
- ☐ None of the above.

(f) (2 pt) What are the disadvantages of disabling interrupts to serialize access to a critical section?

☐ This technique does not work with a multiprocessor.

☐ Disabling interrupts will turn off the hardware timer.

☐ The computer can miss critical events when interrupts are disabled.

☐ User code cannot disable interrupts.

☐ None of the above.

(g) (2 pt) Which of the following MUST occur during a context switch between two threads of different processes?

☐ Computational registers are saved in user memory.

☐ The user stack is saved to the kernel stack.

☐ The file descriptor table is saved to the kernel heap.

☐ Buffered data in memory is flushed to disk.

☐ None of the above.

(h) (2 pt) Which of the following are a result of the "everything is a file" paradigm?

☐ All kernel data structures are stored as files.

☐ User processes interact with I/O devices as if they are files.

☐ Everything stored on disk is part of a file.

☐ One can read and write to a pipe using file syscalls.

☐ None of the above.

(i) (2 pt) Which of the following scheduling algorithms may cause starvation?

☐ Linux CFS

☐ FCFS

☐ Round robin

☐ MLFQS

☐ None of the above.

3. (21 points) Short Answer

Please answer in **THREE SENTENCES OR LESS**. Longer explanations may get no credit.

(a) (3 pt) In Pintos, what does `_start` do besides setting up arguments and calling `main`?

(b) (3 pt) Describe one important difference between `read` and `fread`. How are they related?

(c) (3 pt) Explain the steps needed to establish a pipe between a parent and a child process.

(d) (3 pt) Describe a key difference between Hoare and Mesa semantics.

(e) (3 pt) What role does the Interrupt Vector Table play in protecting the kernel?

(f) (3 pt) Why might the lottery scheduler be a bad choice for real-time scheduling?

(g) **(3 pt)** Rahul plays StarCraft (a computer game). He discovers that inserting `printf` statements into StarCraft code makes StarCraft run faster on his computer. Explain why this happens.

4. (18 points) Parallel File Copy

For online exam takers: you may find it useful to read this problem in a separate PDF instead of scrolling back and forth.

You are making a backup copy of your entire hard drive, but you find that copying a single file at a time takes too long. You decide to put your CS 162 knowledge to good use, and copy multiple files in parallel.

Your program takes in a non-negative integer `n` and two lists of strings, both of length `n`. The first list, `in_files`, contains a list of source files; the second list, `out_files`, contains the corresponding destination file names. Write a program to copy each source file to the corresponding destination. You must create a separate **process** to copy each file. The files should all be copied concurrently.

The main (parent) process must wait for all the files to be copied, and all child processes to exit, before exiting.

Fill in the blanks below to implement your parallel file copy program. You must use the **low-level I/O API** shown in the reference sheet. You may assume calls to `read` and `write` never error. Do **NOT** assume they read/write the full amount.

NOTE: You should only use AT MOST one line per blank. Extra lines will not be graded. Your solution does not need to use all the blanks. You may use loops, conditional statements, and other control structures in your solution as you see fit.

```
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <sys/types.h>
#include <sys/wait.h>
#include <fcntl.h>

#define BUFSIZE 2048

void copy_file(char* src, char* dst) {
  char buf[BUFSIZE];

  int srcfd = open(src, O_RDONLY);
  int dstfd = open(dst, O_WRONLY | O_CREAT | O_TRUNC, 0644);

  int chunk;
  int total;

  while (______________[A]_______________) {
    ______________[B]_______________
    ______________[B]_______________
    ______________[B]_______________
    ______________[B]_______________
    ______________[B]_______________
    ______________[B]_______________
  }

  close(srcfd);
  close(dstfd);
}

void copy(int n, char** in_files, char** out_files) {
  pid_t children[n];
  int i;

  for (i = 0; i < n; i++) {
```

```
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
  }

  ______________[D]_______________
  ______________[D]_______________
  ______________[D]_______________
  ______________[D]_______________
  ______________[D]_______________
}

int main(int argc, char** argv) {
  // Example usage:
  int n = 2;
  char* in_files[] = { "test1.txt", "test2.txt" };
  char* out_files[] = { "copy1.txt", "copy2.txt" };
  // test1.txt should be copied to copy1.txt
  // test2.txt should be copied to copy2.txt
  copy(n, in, out);
  return 0;
}
```

(a) (2 pt) [A] (max 1 line)

(b) (6 pt) [B] (max 6 lines)

(c) (7 pt) [C] (max 11 lines)

(d) (3 pt) [D] (max 5 lines)

5. (25 points) Channels

For online exam takers: you may find it useful to read this problem in a separate PDF instead of scrolling back and forth.

We've seen how locks can be used to synchronize access to shared data. Channels are an alternative model to synchronization. When using channels, one or more threads produce data that is sent through the channel. Other threads then read the data that was sent to the channel.

A buffered channel is one that maintains an internal buffer of a given size. As long as the buffer has space, calls to send should complete quickly (they should just put the new value in the buffer, possibly acquiring some locks if necessary). If the buffer is full, send should wait until space becomes available, unless the caller indicates they do not want to wait (in which case the send operation should do nothing). The `recv` method reads values from the channel in the order in which they were put in the buffer.

In this problem, you will implement a buffered channel that transports integers between multiple threads in the same process. This is the interface you must implement:

```
#include <stdlib.h>
#include <stdbool.h>
#include <pthread.h>

typedef struct chan {
  /* TODO: add fields here. */
} chan_t;

// Initializes a channel with buffer size `size`.
void chan_init(chan_t* chan, size_t size);

// Sends a value through the channel.
// If wait is true, wait for space in the channel if necessary.
// If wait is false, no waiting is permitted. It is OK to acquire
// locks, but you must not wait for some other thread to call send/recv.
// Returns true if the send was successful; false otherwise.
bool send(chan_t* chan, int value, bool wait);

// Receive a value from the channel, storing the result in *dst.
// If wait is true, wait for a value to become available.
// If wait is false, no waiting is permitted. It is OK to acquire
// locks, but you must not wait for some other thread to call send/recv.
// Returns true if the receive was successful; false otherwise.
bool recv(chan_t* chan, int* dst, bool wait);

// Frees all memory allocated by chan_init.
// You can assume that this will only be called by a single thread,
// and that no threads will be waiting due to calls to send/recv.
void chan_free(chan_t* chan);
```

Please be sure to read the descriptions of each method, as they contain important details.

Your implementation should be thread-safe. That is, it should behave correctly regardless of the number of threads concurrently calling `send` and `recv`. You can assume that `chan_init` is called and completes before any concurrent use of the channel.

To make your implementation simpler, you may assume you have access to a circular buffer that stores integers. The interface presented by the buffer is shown below:

```
// A circular buffer of fixed size.
//
// The provided methods are NOT thread-safe.
```

```
typedef struct circ_buf {
  /* FIELDS OMITTED */
} circ_buf_t;


// Initializes a circular buffer of the given size.
void circ_buf_init(circ_buf_t* buf, size_t size);

// Puts a value into the buffer.
// Panics if the buffer is already full.
void circ_buf_put(circ_buf_t* buf, int value);

// Gets and removes a value from the buffer.
// Panics if the buffer is already empty.
int circ_buf_get(circ_buf_t* buf);

// Returns the number of entries in the buffer.
size_t circ_buf_len(circ_buf_t* buf);

// Returns the capacity of the buffer.
// Will always equal the size used to initialize the buffer.
size_t circ_buf_capacity(circ_buf_t* buf);

// Frees all resources allocated by `circ_buf_init`.
// You must not use the buffer after calling this method.
void circ_buf_free(circ_buf_t* buf);
```

The buffer implementation is NOT thread-safe; it is your job to synchronize access to the buffer.

Using the provided buffer API, as well as `pthread` locks, semaphores, and/or condition variables, implement the required channel methods by filling in the blanks below. You may find the reference sheet useful.

NOTE: You should only use AT MOST one line per blank. Extra lines will not be graded. Your solution does not need to use all the blanks. You may use loops, conditional statements, and other control structures in your solution as you see fit.

```
#include "channel.h"
#include <stdio.h>

typedef struct chan {
    circ_buf_t buf;
    ______________[A]_______________
    ______________[A]_______________
    ______________[A]_______________
    ______________[A]_______________
    ______________[A]_______________
    ______________[A]_______________
    ______________[A]_______________
    ______________[A]_______________
    ______________[A]_______________
    ______________[A]_______________
} chan_t;

void chan_init(chan_t* chan, size_t size) {
    ______________[B]_______________
    ______________[B]_______________
    ______________[B]_______________
```

```
    ______________[B]_______________
    ______________[B]_______________
    ______________[B]_______________
    ______________[B]_______________
    ______________[B]_______________
    ______________[B]_______________
    ______________[B]_______________
}

bool send(chan_t* chan, int value, bool wait) {
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
    ______________[C]_______________
}

bool recv(chan_t* chan, int *dst, bool wait) {
    ______________[D]_______________
    ______________[D]_______________
    ______________[D]_______________
    ______________[D]_______________
    ______________[D]_______________
    ______________[D]_______________
    ______________[D]_______________
    ______________[D]_______________
    ______________[D]_______________
    ______________[D]_______________
    ______________[D]_______________
    ______________[D]_______________
    ______________[D]_______________
    ______________[D]_______________
    ______________[D]_______________
    ______________[D]_______________
    ______________[D]_______________
    ______________[D]_______________
    ______________[D]_______________
    ______________[D]_______________
}
```

```
void chan_free(chan_t* chan) {
    ______________[E]_______________
    ______________[E]_______________
    ______________[E]_______________
    ______________[E]_______________
    ______________[E]_______________
}
```

(a) **(5 pt)** [A] (max **10** lines)

(b) (3 pt) [B] (max **10** lines)

(c) (8 pt) [C] (max **20** lines)

(d) (8 pt) [D] (max **20** lines)

(e) (1 pt) [E] (max **5** lines)

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 *mutex, const pthread_mutexattr_t *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); // up
int sem_wait(sem_t *sem); // down
int pthread_cond_init(pthread_cond_t *cond, pthread_condattr_t *cond_attr);
int pthread_cond_signal(pthread_cond_t *cond);
int pthread_cond_broadcast(pthread_cond_t *cond);
int pthread_cond_wait(pthread_cond_t *cond, pthread_mutex_t *mutex);

/*********************************** 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);

/******************************** 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);
off_t lseek(int fd, off_t offset, int whence);
int dup(int oldfd);
int dup2(int oldfd, int newfd);
int pipe(int pipefd[2]);
int close(int fd);
```

No more questions.