

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.

Preliminaries

(a) Full name

(b) Student ID number

(c) Autograder login (e.g. `student123`)

1. (27.0 points)

For the following questions, select whether they are true or false and explain in two sentences or less. Explanations longer than two sentences or lack of explanation will receive no credit.

(a) (3.0 points)

i. A process's address space in kernel space.

True

False

ii. Explain.

(b) (3.0 points)

i. The OS directly schedules processes, not threads.

True

False

ii. Explain.

(c) (3.0 points)

i. Threads in the same process share the same stack.

True

False

ii. Explain.

(d) (3.0 points)

i. After forking, open files are preserved, so closing a file descriptor in the child process would make the parent process unable to use its own file descriptor that refers to the same file object.

- True
- False

ii. Explain.

(e) (3.0 points)

i. Each user thread has its own separate kernel stack.

True

False

ii. Explain.

(f) (3.0 points)

i. Programs which use multiple threads are always faster than programs that do not.

True

False

ii. Explain.

(g) (3.0 points)

i. During a system call, the kernel validates arguments on the user stack then copies them over to the kernel stack.

True

False

ii. Explain.

(h) (3.0 points)

i. x86 calling convention as talked about in Discussion 0 requires the base pointer to be pushed onto the stack by the callee.

- True
- False

ii. Explain.

(i) (3.0 points)

i. The Linux `exec` system call creates a new process.

True

False

ii. Explain.

2. (18.0 points) Select All

- (a) (3.0 pt) Which of the following that are true regarding I/O?
- High level I/O deals with `FILE*` struct.
 - All I/O related syscalls are defined to all forms of I/O under the “everything is a file” uniformity idea from POSIX.
 - There is one file descriptor table per thread.
 - Low-level I/O is buffered in user space.
 - None of the above.
- (b) (3.0 pt) Which of the following are true regarding x86 calling convention as talked about in Discussion 0?
- `addl (%eax), %ebx` adds the longword in memory at the address `eax` to the longword in `ebx`, and stores the result in `eax`.
 - When the callee returns, the return address is gone but the arguments are still on the stack.
 - Stack alignment means the stack is 16 byte aligned before we push arguments onto the stack.
 - The saved `ebp` memory address is lower than the argument’s.
 - The arguments are pushed onto the stack in reverse order.
- (c) (3.0 pt) Which of the following are true about synchronization primitives?
- Semaphores can be used to implement monitors.
 - Monitors can be used to implement semaphores.
 - Semaphores can be used to achieve mutual exclusion but not scheduling constraints.
 - Calls to signal a condition variable can always be replaced with broadcast when using Hoare semantics assuming behavior of Hoare semantics broadcast is the same as the behavior of Mesa semantics broadcast.
 - None of the above.
- (d) (3.0 pt) Which of the following are true about dual mode operation?
- A thread can execute entirely in user mode.
 - A thread can execute entirely in kernel mode.
 - Starting a new process is the only way to switch from kernel to user mode.
 - Page table entries cannot be modified in user mode.
 - None of the above.
- (e) (3.0 pt) Which of the following must be saved during a context-switch?
- Page table
 - Base pointer
 - Current CPU privilege mode
 - Ready list
 - Program counter

(f) (3.0 pt) Which of the following that are true about interrupts?

- The hardware saves the values for the stack pointer and program counter before jumping to the interrupt handler.
- Interrupts are synchronous events independent of the process.
- Disabling interrupts is guaranteed to implement mutual exclusion.
- Lower priority interrupts may be put on hold by a higher priority interrupt.
- None of the above

3. Short Answer

(a) (12.0 points)

Examine the following function `copy_parts(const char *src, const char *dest)`, which copies certain bytes from the `src` file into the `dest` file.

```
void copy_parts(const char* src, const char* dest) {
    char* buffer = malloc(1000 * sizeof(char));

    int read_fd = open(src, O_RDONLY);
    lseek(read_fd, 500, SEEK_SET);
    read(read_fd, buffer, 150 * sizeof(char));

    int write_fd = open(dest, O_WRONLY);
    lseek(write_fd, 200, SEEK_SET);
    write(write_fd, buffer, 100 * sizeof(char));

    /* breakpoint */

    read(read_fd, buffer, 100 * sizeof(char));
    write(write_fd, buffer + 50, 100 * sizeof(char));

    close(read_fd);
    close(write_fd);
}
```

- i. (4.0 pt) Is the following function's behavior deterministic? If not, what is the issue and how do we fix it?

- ii. (4.0 pt) If there was an issue with the function, suppose we have addressed the issue and the function behavior is now guaranteed to be deterministic. What is the current file position of the `src` and `dest` files at the breakpoint? Explain.

- iii. (4.0 pt) If there was an issue with the function, suppose we have addressed the issue and the function behavior is now guaranteed to be deterministic. Suppose a file `src.txt` has 500 `a`'s followed by 200 `b`'s followed by 300 `c`'s followed by EOF. Suppose a file `dest.txt` has 100 `x`'s followed by EOF. What is the result of calling `copy_parts("src.txt", "dest.txt")`? Explain. You may use up to ten sentences for this question to explain if necessary.



(b) (4.0 pt) Can a user thread acquire a lock without entering the kernel? Explain.

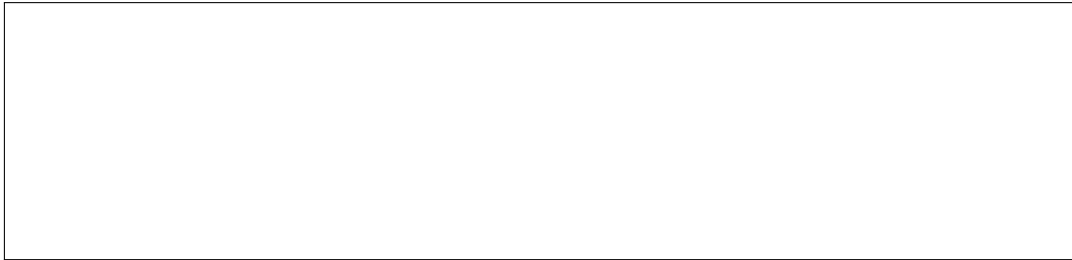
(c) (4.0 pt) Explain the purpose of system calls, and why they cannot be executed in user mode.

(d) (4.0 pt) Explain two reasons why high-level I/O uses a user space buffer.

(e) (4.0 pt) What happens when a multithreaded program calls fork?

(f) (4.0 pt) In Pintos, is there a need to switch page directory base register when switching from user to kernel mode? Why or why not?

(g) (4.0 pt) What does the `$0x30` represent in `int $0x30` in Pintos?



4. (18.0 points) Shared Data Revisited

The following is a modified version of the shared data question from Discussion 2, and is related to the implementation of the `wait` syscall in Project 1. For this problem, we're going to implement a wait function that allows one thread (the main thread) to wait for another thread to finish writing some data.

We're going to assume we don't have access to `pthread_join` for this problem. Instead, we're going to use synchronization primitives (locks and semaphores). We need to design a struct for sharing information between the two threads.

We also need to implement three functions to initialize the shared struct and synchronize the 2 threads. `initialize_shared_data` will initialize our shared struct. `wait_for_data` (called by the main thread) will block until the data is available. `save_data` (called by the child thread) will write 162 to the struct. Another requirement is that the shared data needs to be freed once as soon as it is no longer in use. When using synchronization primitives, the critical section should be as few lines as possible while still ensuring a race free environment.

Fill in the blanks so that main will always print "Parent: Data is 162". You may assume the program compiles without error and all library calls will succeed.

```
typedef struct shared_data {
    sem_t sem;
    int data;
    -----
} shared_data_t;

void initialize_shared_data(shared_data_t* shared_data) {
    sem_init(&shared_data->sem, 0, 0);
    shared_data->data = -1;
    -----
}

int wait_for_data(shared_data_t* shared_data) {
    sem_wait(&shared_data->sem);
    int data = shared_data->data;
    -----
    return data;
}

void* save_data(void* shared_pg) {
    shared_data_t* shared_data = (shared_data_t*)shared_pg;
    shared_data->data = 162;
    sem_post(&shared_data->sem);
    -----
    return NULL;
}

int main() {
    shared_data_t* shared_data = malloc(sizeof(shared_data_t));
    initialize_shared_data(shared_data);
    pthread_t tid;
    int error = pthread_create(&tid, NULL, &save_data, (void*)shared_data);
    int data = wait_for_data(shared_data);
    printf("Parent : Data is % d\n", data);
    -----
    return 0;
}
```

(a) (3.0 pt) Fill in the missing code (if any) for `shared_data_t`. You may use up to five lines.

(b) (3.0 pt) Fill in the missing code (if any) for `initialize_shared_data`. You may use up to five lines.

(c) (4.0 pt) Fill in the missing code (if any) for `wait_for_data`. You may use up to 10 lines.

A large, empty rectangular box with a thin black border, intended for the student to write their code for the `wait_for_data` function. The box is currently blank.

(d) (2.0 pt) Fill in the missing code (if any) for `save_data`. You may use up to 10 lines.

A large empty rectangular box with a thin black border, intended for the student to write their code for the `save_data` function. The box is currently blank.

(e) (2.0 pt) Fill in the missing code (if any) for `main`. You may use up to 10 lines.

(f) (4.0 pt) What kind of synchronization, if any, is needed for `shared_data->data` in both `wait_for_data` and `save_data`? Explain.

5. (22.0 points) Santa's List

Santa is digitizing his accounting department and would like to track naughty and nice children on his computer. He has a file for each person in the form of `childX.txt` (e.g. `child0.txt`) that contains string records representing the good and bad deeds someone has done all year. For example,

```
child0.txt
```

```
GOOD attended live lecture and asks questions.
GOOD finished assignments in a timely manner.
BAD didn't study for midterm.
```

```
child1.txt
```

```
BAD didn't attend any discussion.
BAD didn't help his teammates during projects.
```

```
child2.txt
```

```
GOOD answered a lot of questions on Piazza.
BAD didn't fill out course evaluations.
```

Fill in the code below that determines whether a child has performed more good deeds than bad deeds. As there are many children, create a new thread for each file for concurrency. When all threads are done counting all files, the program should write the filenames with positive karma to `nice.txt` and others to `naughty.txt`. Separate each filename by a newline with a trailing newline at the end of the file.

```
typedef struct thread_data {
    int tid;
    pthread_t thread;
    char filename[128];
    int karma;
} thread_data_t;

void* count_good_bad(void* threadarg) {
    ----- [A] -----;
    ----- [B] -----;
    char* line = NULL;
    size_t len = 0;
    ssize_t read;
    while ((read = getline(&line, &len, fp)) != -1) {
        if (strncmp(line, "GOOD", 4) == 0) {
            tdata->karma++;
        } else if (strncmp(line, "BAD", 3) == 0) {
            tdata->karma--;
        }
    }
    pthread_exit(NULL);
}

int main(int argc, char* argv[]) {
    int num_children = atoi(argv[1]);
    thread_data_t thread_data_array[num_children];

    for (int i = 0; i < num_children; i++) {
        thread_data_array[i].tid = i;
        /* Stores child<i>.txt (e.g. child0 for i = 0) in to thread_data_array[i].filename */
        sprintf(thread_data_array[i].filename, "child%d.txt", i);
        thread_data_array[i].karma = 0;
    }
}
```

```
    -----[C]-----;
}

for (int i = 0; i < num_children; i++) {
    pthread_join(thread_data_array[i].thread, NULL);
}

FILE* fp_nice = fopen("nice.txt", "w");
FILE* fp_naughty = fopen("naughty.txt", "w");

for (-----[D]-----) {
    if (-----[E]-----) {
        -----[F]-----;
    } else {
        -----[G]-----;
    }
}

fclose(fp_nice);
fclose(fp_naughty);

return 0;
}
```

With the above example, the files would look like

nice.txt

child0.txt

child2.txt

naughty.txt

child1.txt

(a) (2.0 pt) [A]

(b) (2.0 pt) [B]

(c) (2.0 pt) [C]

(d) (2.0 pt) [D]

(e) (2.0 pt) [E]

(f) (2.0 pt) [F]

(g) (2.0 pt) [G]

(h) (4.0 pt) What would happen if we did not call `pthread_join`?

(i) (4.0 pt) How would you implement this using processes rather than threads?

6. Improving pwords

Recall `pwords` from Homework 1 which used multiple threads to count words. Specifically, each thread handles one file and writes to a shared `word_count_list_t` struct that is protected by a lock. As you noticed, its performance does not improve much over the single-threaded implementation. In this problem, we will attempt to remedy this situation.

In the following code, we are still using N (assume at least 2) “worker threads”, one to process each input file, but each thread will write to its own (non-shared) `word_count_list_t` struct first. Then, we are going to use one separate “merger thread” to merge any 2 ready `word_count_list_t` structs. All these `word_count_list_t` structs will be communicated through a currently **non-thread safe** `queue_t` struct.

The skeleton of the non-thread safe program is shown below. Your task is to make it thread-safe by **using the two provided semaphores only**. Your code may only include semaphore operations `sem_wait` and `sem_post`. Your solution must maximize the amount of work that can execute in parallel.

You may assume all calls to `fopen`, `fclose`, `pthread_create`, `pthread_join`, `malloc`, and other provided functions succeed.

```

/* FIFO queue. */
typedef struct queue queue_t;
/* Initializes empty queue into QUEUE. */
void queue_init(queue_t* queue);
/* Pops the front from QUEUE returns it. Undefined behavior when queue is empty. */
wordcount_list_t* queue_pop(queue_t* queue);
/* Pushes WCLIST to the end of QUEUE. */
void queue_push(queue_t* queue, word_count_list_t* wclist);

/* Initializes an empty word count list into WCLIST. */
void words_init(word_count_list_t* wclist);
/* Counts words into WCLIST from INFILE. */
void words_count(word_count_list_t* wclist, FILE* infile);
/* Sorts contents of WCLIST descending order in-place. */
void words_sort(word_count_list_t* wclist);
/* Prints contents of WCLIST. */
void words_print(word_count_list_t* wclist);
/* Returns a new merged word count list that combines the word counts of WCLIST1 and WCLIST2. */
word_count_list_t* words_merge(word_count_list_t* wclist1, word_count_list_t* wclist2);

int N;
sem_t sem_mutex;
sem_t sem_items;
queue_t queue;

typedef struct worker_args {
    int tid;
    char* name;
    pthread_t thread;
    word_count_list_t* wclist;
} worker_args_t;

void* worker_thread_func(void* arg) {
    worker_args_t* worker_args = (worker_args_t*)arg;

    FILE* infile = fopen(worker_args->name, "r");
    words_count(worker_args->wclist, infile);
    fclose(infile);

```

```

    -----[A]-----;
    queue_push(&queue, wclist);
    -----[B]-----;
    -----[C]-----;
}

void* merger_thread_func(void* arg) {
    for (int i = 0; i < N - 1; i++) {
        -----[D]-----;
        -----[E]-----;
        -----[F]-----;
        word_count_list_t* wclist1 = queue_pop(&queue);
        word_count_list_t* wclist2 = queue_pop(&queue);
        -----[G]-----;

        word_count_list_t* merged = words_merge(wclist1, wclist2);

        -----[H]-----;
        queue_push(&queue, merged);
        -----[I]-----;
        -----[J]-----;
    }
}

int main(int argc, char* argv[]) {
    N = argc - 1;
    sem_init(&sem_mutex, 0, 1);
    sem_init(&sem_items, 0, 0);
    queue_init(&queue);

    worker_args workers[N];
    for (int i = 0; i < N; i++) {
        workers[i].name = argv[i + 1];
        workers[i].tid = i;
        workers[i].wclist = malloc(sizeof(word_count_t));
        words_init(&workers[i].wclist);
        pthread_create(&workers[i].thread, NULL, worker_thread_func, &workers[i]);
    }

    pthread_t merger_thread;
    pthread_create(&merger_thread, NULL, merger_thread_func, NULL);

    for (int i = 0; i < N; i++) {
        pthread_join(workers[i].thread, NULL);
    }

    pthread_join(merger_thread, NULL);

    word_count_list_t* final_result = queue_pop(&queue);
    words_sort(final_result);
    words_print(final_result);

    return 0;
}

```

(a) (3.0 pt) [A]

(b) (3.0 pt) [B]

(c) (3.0 pt) [C]

(d) (3.0 pt) [D]

(e) (3.0 pt) [E]

(f) (3.0 pt) [F]

(g) (3.0 pt) [G]

(h) (3.0 pt) [H]

(i) (3.0 pt) [I]

(j) (3.0 pt) [J]

(k) (3.0 pt) Describe the purpose of the two semaphores `sem_mutex` and `sem_items`.

(l) (4.0 pt) Would launching 2 merger threads work? Why or why not?

(m) (4.0 pt) Can this approach be faster than original `pwords`? Why or why not?

7. Multithreading Fun

Consider the following two threads. Assume each instruction is atomic.

```
# =====  
# Syntax  
# load r, a    loads a into r  
# store r, b   stores r into b  
# add  r, c    adds c to r and stores into r  
# sub  r, c    subtracts c from r and stores into r  
# mul  r, c    multiples r by c and stores into r  
# =====
```

```
x = 1  
y = 3
```

```
# Thread A  
load r1, x  
add  r1, 1  
store r1, y
```

```
# Thread B  
load r1, y  
sub  r1, 1  
store r1, x
```

- (a) (4.0 pt) What are all the possible values of x and y? List them as tuples separated by commas (e.g. (1, 2), (3, 4))

- (b) (4.0 pt) Now consider a third thread executing in parallel with threads A and B. Is there any sequence of execution that produces $x = 4$ and $y = 4$ after all threads complete? If yes, provide an example of such execution by writing the exact execution order where each line is the line of code and which thread it belongs to. If not, explain why not.

```
# Thread C
load r1, x
mul r1, 2
store r1, x
```

- (c) (4.0 pt) How would you use locks in thread A and thread B such that these threads can produce no more than three distinct values between both x and y ?

8. References

```

/* Semaphore */
int sem_init(sem_t *sem, int pshared, unsigned int value);
int sem_wait(sem_t *sem);
int sem_post(sem_t *sem);

/* Lock */
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);

/* Condition Variable */
int pthread_cond_init(pthread_cond_t *cond, const pthread_condattr_t *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);

/* Process */
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[]);

/* Thread */
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);
void pthread_exit(void *retval);

/* High-Level IO */
FILE *fopen(const char *pathname, const char *mode);
int fclose(FILE *stream);
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);
fprintf(FILE * restrict stream, const char * restrict format, ...);

/* Low-Level IO */
int open(const char *pathname, int flags);
int close(int fd);
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);

/* Socket */
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);

```


9. Clarifications

The following are selected clarifications made during the exam. Some clarifications have not been included and instead corrected within the exam since they were typos and didn't significantly affect the exam.

True or False

- Regarding Pintos for (5)

Select All

- “Defined to” means “applicable to” for (1)
- Thread execution does not include creation or exit for (4)
- Switching refers to thread switching for (5)
- “Select all of the following that are true of interrupts.” is the prompt for (6)

Short Answer

- Assume all syscalls don't error when called for (1).

Santa's List

- Nonnegative karma (karma greater than or equal to 0) should be written to `nice.txt`. everything else will go to `naughty.txt`.
- `ssize_t getline(char ** restrict linep, size_t * restrict linecapp, FILE * restrict stream)` is the method header for `getline`

Multithreading Fun

- `r1` is a global variable.
- Assume all threads run to completion.
- Thread C is not used for (c).
- Describe a situation with exactly three distinct values for (c).

No more questions.