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. There are a total of 75 points in this exam. You have 75 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) Discussion TA's Full Name

(d) 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-sheets 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. **(14 points)   True/False**

   Please explain your answer in **TWO SENTENCES OR LESS**. Longer explanations will **GET NO CREDIT**. Answers without an explanation or that just rephrase the question will **GET NO CREDIT**.

   (a) **(2 points)**

       i. If you call a function `f(a, b, c)` using the x86 calling convention, the resulting stack would look like the following:

Stack Diagram
padding
a
b
c
return address

       Note: In the diagram above, addresses increase upward.

       ○ True.

       ○ False.

       ii. Explain.

   (b) **(2 points)**

       i. The `pipe()` system call creates a file on disk and returns file descriptors to read and write to it.

       ○ True.

       ○ False.

       ii. Explain.

(c) (2 points)

    i. In PintOS, we validate syscall arguments in user memory.

       ○ True.

       ○ False.

    ii. Explain.

(d) (2 points)

    i. In Unix, while `fork()` creates a new process and preserves the parent process, `exec()` creates a new process and kills the parent process.

       ○ True.

       ○ False.

    ii. Explain.

(e) **(2 points)**

    i. Two processes can agree to share memory by both reading and writing to the virtual address `0xFFFF0000`.

      ○ True.

      ○ False.

    ii. Explain.

(f) **(2 points)**

    i. If a single process opens a file twice, a new file descriptor is created for each **open** call and the two file descriptors both point to the different file descriptions.

      ○ True.

      ○ False.

    ii. Explain.

(g) (2 points)

   i. If you fork a process, changes made to the forked file descriptions are reflected in both the parent and child processes.

     ○ True.

     ○ False.

  ii. Explain.

2. (14 points)   Multiple Select

(a) **(2 pt)** Select the following statements about sockets that are true.

☐ Sockets provide bidirectional communication between processes.

☐ Open socket file descriptors can be modified using the `lseek()` syscall.

☐ A single connection socket can be used to talk to multiple (host, port) endpoints.

☐ TCP sockets guarantee reliable, in-order communication.

☐ None of the above.

(b) **(2 pt)** Select the following statements about the roles of the operating system that are true.

☐ The operating system provides common services such as file systems.

☐ CPU time is an example of a shared resource the operating system manages.

☐ A process is not allowed to modify the address space of another process.

☐ A user program can directly access a file without making a syscall.

☐ None of the above.

(c) **(2 pt)** Select the following statements about x86 calling convention that are true.

☐ The caller must always save caller-saved general purpose registers.

☐ In the callee epilogue, the callee can deallocate local variables by setting the `ebp` to the `esp`.

☐ After the instruction `pushl %ebx`, the `esp` will be decremented by 4 bytes.

☐ After pushing the parameters onto the stack, the stack is 16-byte aligned.

☐ None of the above.

(d) **(2 pt)** Select the following statements about syscalls in PintOS that are true.

☐ Syscalls by user programs go directly to the corresponding function implementation.

☐ The `exit` syscall is eventually called at the end of a user program.

☐ Syscalls are handled in kernel mode.

☐ Syscalls are handled asynchronously.

☐ None of the above.

(e) **(2 pt)** Select the following statements about hardware support for locks that are true.

☐ Atomic instructions beyond load and store are required to make a correctly functioning method of mutual exclusion for a uniprocessor.

☐ Busy-waiting on a uniprocessor is generally as efficient as implementing locks by disabling and enabling interrupts.

☐ A lock implemented using `test_and_set(&lock_val)` as described in lecture can have `lock_val` initialized to any value.

☐ All correct lock implementations for uniprocessors function correctly for multiprocessors.

☐ None of the above.

(f) (2 pt) In class, we discussed how the `fork()` creates a copy of the old process. What values/structures would `fork` copy into the new process? (Note that in this question, you can assume that we are not using the copy-on-write optimization.)

☐ Code Segment

☐ PID

☐ ESP

☐ EIP

☐ None of the above.

(g) (2 pt) Select the following statements about signals that are true.

☐ The `kill` syscall can only be used to send the `SIGKILL` signal.

☐ Programs can override the handlers for any signal using the `sigaction` function.

☐ The operating system needs to transition from user mode into kernel mode in order to allow the current process to handle a signal.

☐ The kernel indicates to a process that a segmentation fault has occurred by sending a `SIGSEGV` signal.

☐ None of the above.

3. (8 points)   Short Answer

(a) **(2 pt)** Imagine we are trying to write out to a file from a buffer. Name 2 differences that we need to make in our implementation if we use the low-level `write()` compared to if we use the high-level `fwrite()`, and why you need to make them.

(b) **(2 points)**

Recall the different mechanisms for mode transfer that we learned about in class:

- Syscalls
- Interrupts
- Exceptions

For each of the following, list which of the above 3 mode transfer mechanisms most closely matches the described situation. You may use an answer more than once.

i. **(0.5 pt)** A timer fires, which indicates that a different process will now be scheduled.

ii. **(0.5 pt)** The user programs attempts to access illegal memory and a segfault occurs.

iii. **(0.5 pt)** An I/O operation which writes to disk while other work is done on the CPU.

iv. (0.5 pt) User program invokes a function to write to a file.


(c) (2 points)

Indicate the output produced by running the following code. Please separate the output in terms of what would be seen in the terminal and what would be seen in the file. *Assume that the printf statements are atomic.*

```c
void* helper(void* args) {
    int *fd = (int*) args;
    printf("Printing from thread helper\n");
    dup2(*fd, STDOUT_FILENO);
    return NULL;
}

int main() {
    printf("Starting main\n");
    int file_fd = open("test.txt", O_WRONLY | O_TRUNC | O_CREAT, 0666);
    int* temp_fd = (int*) malloc(sizeof(int));
    *temp_fd = dup(STDOUT_FILENO);
    pthread_t thread;
    dup2(file_fd, STDOUT_FILENO);
    pthread_create(&thread, NULL, &helper, temp_fd);
    pthread_join(thread, NULL);
    printf("Ending main\n");
    return 0;
}
```

i. In terminal:

ii. In `file.txt`:

(d) (2 points)

Suppose you were given the following implementation of a producer-consumer queue, based on the code provided in the lecture slides. Note that `dequeue` will segfault if the `queue` that is passed in is empty.

In `queue.h`:

```
typedef struct queue {
    // ... omitted ...
} queue_t;

void enqueue(queue_t* queue, int val);
// dequeue segfaults if the queue is empty
int dequeue(queue_t* queue);
bool is_empty(queue_t* queue);
```

In `pc_queue.c`:

```
#include <pthread.h>
#include "queue.h"

typedef struct pc_queue {
    pthread_mutex_t lock;
    pthread_cond_t cv;
    queue_t queue;
} pc_queue_t;

void produce(pc_queue_t* pcq, int item) {
    pthread_mutex_lock(&pcq->lock);
    enqueue(&pcq->queue, item);
    pthread_cond_signal(&pcq->cv);
    pthread_mutex_unlock(&pcq->lock);
}

int consume(pc_queue_t* pcq) {
    pthread_mutex_lock(&pcq->lock);
    _____ {
        pthread_cond_wait(&pcq->cv, &pcq->lock);
    }
    int item = dequeue(&pcq->queue);
    pthread_mutex_unlock(&pcq->lock);
    return item;
}
```

    i. **(0.5 pt)** Given that the `pthread` library uses Mesa semantics, what is the best line to fill in the blank above?

ii. **(0.5 pt)** Suppose instead that the `pthread` library used Hoare semantics. How could you simplify your answer to the previous question?

iii. **(1 pt)** Suppose that the programmer tried to fill in the blank line assuming that `pthread` used Hoare semantics, while the actual library uses Mesa semantics. Would any bugs result? If so, what might happen? If not, explain why.

4. (10 points)   My Very Own fread!

Provide an implementation of `my_fread` using the low-level file I/O API. `my_fread` should behave similarly to `fread`, where we read data using the low-level file API and buffer the results to speed up future calls.

You may not use `fread` or any of the other functions defined in `stdio.h`. You may assume that the number of bytes requested by `my_fread` will be at most the number of bytes remaining in the open file.

```
#define BUF_SIZE 1024
#define MIN(X, Y) (((X) < (Y)) ? (X) : (Y))

typedef struct file {
    int fd;
    char buffer[BUF_SIZE];
    int buf_offset;         // initialized to 0
    bool has_data;          // initialized to false
} File;

int refill_buffer(File* file) {
    ssize_t bytes_read;
    int read_so_far = 0;
    while ((bytes_read = _____[A]_____) > 0) {
        read_so_far += _____[B]_____;
    }
    file->buf_offset = _____[C]_____;
    file->has_data = true;
    return read_so_far;
}

/* You may assume that `count` is at most the remaining
 * number of bytes in the open file `file->fd`. */
void my_fread(char* dest, size_t count, File* file) {
    if (!file->has_data) {
        _____[D]_____;
    }
    int written_so_far = 0;
    while (written_so_far < count) {
        int bytes_to_copy = MIN(_____[E]_____);
        memcpy(_____[F]_____);
        file->buf_offset += _____[G]_____;
        if (file->buf_offset == BUF_SIZE) {
            _____[H]_____;
        }
        written_so_far += _____[I]_____;
    }
}
```

(a) **(2 pt)** [A] (max 1 expression)

(b) **(1 pt)** [B] (max 1 expression)

(c) **(1 pt)** [C] (max 1 expression)

(d) **(1 pt)** [D] (max 1 line)

(e) **(1 pt)** [E] (parameters)

(f) **(1 pt)** [F] (parameters)

(g) **(1 pt)** [G] (max 1 expression)

(h) **(1 pt)** [H] (max 1 line)

(i) **(1 pt)** [I] (max 1 expression)

5. (15 points)   Reminders

Saahya is creating an interactive productivity tool in PintOS to help her team track tasks in their CS162 project. One feature of this tool is that users can schedule reminder messages to be printed to the screen at regular intervals. All times are specified in seconds.

For example: At `time=100`, a user might schedule a "help!" reminder with interval 10. Then, at `time=110`, `120, 130, 140, ...`, "help!" will be printed to the screen.

Note that Saahya is implementing this user program in a modified version of PintOS which includes the following two syscalls:

```
int time(); // returns current time in seconds
void sleep(unsigned int seconds); // puts the current thread to sleep for X seconds.
```

For all parts, you may assume that reminder strings are under 100 characters, and that all the intervals are on the order of minutes to hours (so you don't need to worry about integer overflow). Currently, the main thread of Saahya's tool looks like the following:

```
/* Returns input string from the user.
 * If there is no input ready, put the calling thread to sleep until input is available. */
char* get_user_input();

int main() {
    while (true) {
        char* input = get_user_input();
        // ... process user input ...
    }
}
```

(a) (1 pt) Saahya intends to implement the reminder tracking/printing logic on the main thread in the process user input section. Why is this a bad idea?

(b) (14 points)

Taking into account your concerns, Saahya refactors her existing codebase in preparation for adding reminders:

```
typedef struct reminder {
    char* message; // The reminder message
    int interval; // The interval at which the message is sent
    int target_time; // The timestamp marking when the reminder should be sent next;
    struct list_elem elem;
} reminder_t;

typedef struct reminder_list {
    struct list list; // List of reminder_t's
    struct lock lock; // Lock for accessing the list
    struct condition cond;
} reminder_list_t;

int main() {
    reminder_list_t reminders;
    initialize(&reminders);

    pthread_t* thread;
    pthread_create(&thread, NULL, runner, &reminders);

    /* ... long-running logic which will occasionally call
        add_reminder based on user input ... */
}

void initialize(reminder_list_t* reminders) {
    // TODO
}

void add_reminder(reminder_list_t* reminders, char* message, int interval) {
    // TODO
}

void wake_after_delay(void* arg) {
    // TODO
```

```
}

void runner(void* arg) {
    // TODO
}
```

Your task is to help Saahya fill in the unimplemented functions.

i. **(1 pt)** Consider the following example:

- At `time=0`, main calls `add_reminder(reminders, ''Work on design!'', 15);`
- At `time=20`, main calls `add_reminder(reminders, ''Stay focused'', 5);`
- At `time=31`, main exits the process, killing all threads.

Please list the `(time, message)` pairs for every message sent.

ii. (1.5 pt) Implement the `initialize` function, which initializes the `reminders` struct.

```
void initialize(reminder_list_t* reminders) {
    _____[A (x3)]_____
}
```

[A] (max 3 lines)

```



```

iii. (3.5 points)

Implement the `add_reminder` function, which should add a new reminder struct into the `reminders` list and then signal the condition variable. Ensure that all necessary synchronization is in place.

Note: You may not assume that `*message` will persist beyond this function call, and you are not responsible for freeing it.

```
void add_reminder(reminder_list_t* reminders, char* message, int interval) {
    reminder_t* reminder = _____[A]_____
    _____[B (x3)]_____

    _____[C]_____
    _____[D]_____
    cond_signal(&reminders->cond, &reminders->lock);
    _____[E]_____
}
```

A. (0.5 pt) [A] (max 1 expression)

```


```

B. (1.5 pt) [B] (max 3 lines)

```




```

C. (0.5 pt) [C] (max 1 line)

```


```

D. (0.5 pt) [D] (max 1 line)

E. (0.5 pt) [E] (max 1 line)

iv. **(3 points)**

In preparation for implementing runner, implement the `wake_after_delay` helper function. This function should sleep until `info->target_wake_time`, then signal the condition variable.

```
typedef struct delay_info {
    struct condition* cond; // Pointer to the reminder_list's cond
    struct lock* lock; // Pointer to the reminder_list's lock
    int target_wake_time; // Target time to signal cond
} delay_info_t;

void wake_after_delay(void* arg) {
    delay_info_t* info = (delay_info_t*) arg;
    int time_till_wake = _____[A]_____ - time();
    if (_____[B]_____) {
        sleep(time_till_wake);
    }

    _____[C]_____;
    cond_signal(_____[D]_____, _____[E]_____);
    _____[F]_____;

    free(info);
}
```

A. (0.5 pt) [A] (max 1 expression)

B. (0.5 pt) [B] (max 1 expression)

C. (0.5 pt) [C] (max 1 line)

D. (0.5 pt) [D] (max 1 expression)

E. (0.5 pt) [E] (max 1 expression)

F. (0.5 pt) [F] (max 1 line)

v. (5 points)

Finally, implement the runner function. This function should implement the reminder-sending feature. If n is the number of notifications sent over some timespan, the number of while loop iterations must be at most $2n$ in all cases.

```
void runner(void* arg) {
    reminder_list_t* reminders = (reminder_list_t*) arg;
    _____[A]_____;
    while (true) {
        cond_wait(&reminders->cond, &reminders->lock);
        int t = time();
        int target_wait_time = INT_MAX; // target_time of the closest reminder

        for (_____[B]_____; _____[C]_____; _____[D]_____) {
            reminder_t* reminder = _____[E]_____;
            if (_____[F]_____) {
                printf(reminder->message);
                _____[G]_____;
            }
            if (_____[H]_____) {
                target_wait_time = reminder->target_time;
            }
        }

        delay_info_t* info = _____[I]_____;
        info->cond = &reminders->cond;
        info->lock = &reminders->lock;
        info->target_wake_time = target_wake_time;
        _____[J]_____;
    }
}
```

A. (0.5 pt) [A] (max 1 line)

```
```

B. (0.5 pt) [B] (max 1 expression)

```
```

C. (0.5 pt) [C] (max 1 expression)

```
```

D. (0.5 pt) [D] (max 1 expression)

```
```

E. (0.5 pt) [E] (max 1 expression)

F. (0.5 pt) [F] (max 1 expression)

G. (0.5 pt) [G] (max 1 line)

H. (0.5 pt) [H] (max 1 expression)

I. (0.5 pt) [I] (max 1 expression)

J. (0.5 pt) [J] (max 1 line)

6. (14 points)   Hybrid Thundering Locks

WebKit—a browser engine developed by Apple—features optimized custom implementations of several synchronization primitives. In particular, it includes a high-throughput "adaptive lock" implementation which performs especially well under microcontention. *Microcontended locks* are locks that protect short critical sections, so they should be available soon after failed attempts to acquire them.

In this problem, you will implement a similar "hybrid thundering lock" in C by leveraging the simplified `ParkingLot` API and atomic operations. Note that limited busy waiting in parts of this problem is allowed, as long as it remains bounded within the described limits.

You have access to the following `ParkingLot` API functions, which enable efficient and safe thread queueing and dequeuing, similar to `futex`:

Parking threads effectively puts them to sleep and queues them.

```
/* Compares the `lock` and `desired_state` and parks the current thread if they match. */
void compare_and_park(hybrid_lock_t* lock, hybrid_lock_t desired_state);
```

Unparking threads dequeues them and wakes them up.

```
/* Unpark one thread that is waiting on `lock`. */
void unpark_one(hybrid_lock_t* lock);
```

```
/* Unpark all threads that are waiting on `lock`. */
void unpark_all(hybrid_lock_t* lock);
```

Assume you also have access to the compare-and-swap (CAS) and swap atomic operations:

```
/* pseudocode for behavior of the atomic compare and swap instruction */
bool CAS(int* addr, int expr1, int expr2) {
  if (*addr == expr1) {
    *addr = expr2;
    return true;
  } else {
    return false;
  }
}
```

```
/* pseudocode for behavior of the atomic swap instruction */
int swap(int* addr, int expr1) {
  int temp = *addr;
  *addr = expr1;
  return temp;
}
```

(a) (1 points)   Lock Init

Each lock has 3 distinct states, which we represent with the following `typedef` and constants:

```
#define UNLOCKED 0 // lock is free
#define LOCKED 1 // lock is held and no parked threads are waiting for it
#define LOCKED_AND_PARKED 2 // lock is held and at least 1 parked thread is waiting for it

typedef int hybrid_lock_t;
```

Using the provided API's, implement the required lock functions by filling in the blanks below. You may find the reference sheet useful.

```
void lock_init(hybrid_lock_t* lock) {
    _____[A]_____
}
```

i. **(1 pt)** [A] (max 1 line)

(b) (8 points)   Lock Acquire

Lock acquisition will have 2 phases:

   i. A primary bounded spin-lock phase which spins for a maximum of 40 times before moving on
  ii. A secondary slow phase which should park the requesting thread if the lock is not free.

```
void lock_acquire(hybrid_lock_t * lock) {
    // Bounded spin-lock phase
    for (_____[B]_____) {
        if ( _____[C]_____ && CAS(_____[D]_____)) {
            return;
        }

        if ( _____[E]_____) {
            break;
        }
    }


    // Slow phase
    while (true) {
        if ( _____[F]_____ && CAS(_____[G]_____)) {
            return;
        }
        CAS(_____[H]_____);
        compare_and_park(_____[I]_____);
    }
}
```

  i. (1 pt) [B] (max 3 lines)

  ii. (1 pt) [C] (max 1 expression)

  iii. (1 pt) [D] (parameters)

  iv. (1 pt) [E] (max 1 expression)

v. **(1 pt)** [F] (max 1 expression)

viii. **(1 pt)** [G] (parameters)

vii. **(1 pt)** [H] (parameters)

viii. **(1 pt)** [I] (parameters)

(c) (3 points)   Lock Release

Upon lock release, if any parked threads are waiting for the lock, they should be unparked, thus releasing a "thundering herd" of threads that will try to acquire it.

```
void lock_release(hybrid_lock_t * lock) {
    if (CAS(_____[J]_____)) {
        return;
    }
    if (swap(lock, Unlocked) == _____[K]_____) {
        _____[L]_____
    }
}
```
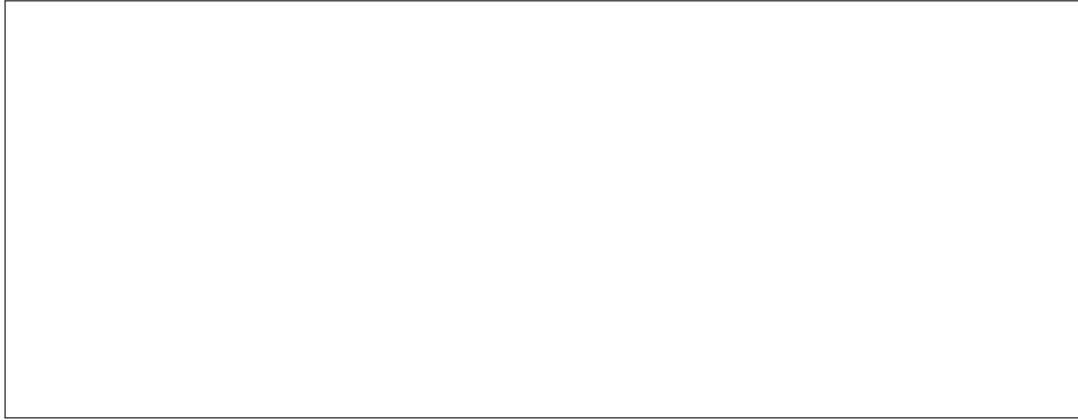
  i. (1 pt) [J] (parameters)

  ii. (1 pt) [K] (max 1 expression)

 iii. (1 pt) [L] (max 1 line)

(d) (2 points)   Design Review

    i. (2 pt) Assume all uses of the hybrid thundering lock are under microcontention as defined above. What is one potential advantage and one potential disadvantage of this particular lock design?

7. Reference Sheet (feel free to remove)

```
/* 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[]);

/* pthreads */
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);

/* pthread Semaphore interface */
int sem_init(sem_t *sem, int pshared, unsigned int value);
int sem_wait(sem_t *sem); /* The p() or down() operation */
int sem_post(sem_t *sem); /* The v() or up() operation */

/* pthread Lock/mutex operations */
pthread_mutex_t mutex = PTHREAD_MUTEX_INITIALIZER;
int pthread_mutex_init(pthread_mutex_t *mutex, const pthread_mutexattr_t *attr);
int pthread_mutex_destroy(pthread_mutex_t *mutex);
int pthread_mutex_lock(pthread_mutex_t *mutex);
int pthread_mutex_trylock(pthread_mutex_t *mutex);
int pthread_mutex_unlock(pthread_mutex_t *mutex);

/* pthread Condition Variable */
pthread_cond_t cond = PTHREAD_COND_INITIALIZER;
int pthread_cond_init(pthread_cond_t *cond, const pthread_condattr_t *attr);
int pthread_cond_destroy(pthread_cond_t *cond);
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);

/* Pintos locks */
void lock_init(struct lock *lock);
void lock_acquire(struct lock *lock);
void lock_release(struct lock *lock);

/* Pintos semaphore interface */
void sema_init(struct semaphore *sema, unsigned value);
void sema_down(struct semaphore *sema);
void sema_up(struct semaphore *sema);

/* Pintos condition variables */
void cond_init(struct condition *cond);
void cond_wait(struct condition *cond, struct lock *lock);
void cond_signal(struct condition *cond, struct lock *lock);
void cond_broadcast(struct condition *cond, struct lock *lock);

/* Pintos Readers/Writers Locks */
void rw_lock_init(struct rw_lock*);
void rw_lock_acquire(struct rw_lock*, bool reader);
void rw_lock_release(struct rw_lock*, bool reader);
```

```
/* Pintos List */
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);


/* Strings */
char *strcpy(char *dest, char *src);
char *strdup(char *src);


/* Interrupt enable/disable */
enum intr_level {};
enum intr_level intr_get_level(void)
enum intr_level intr_set_level(enum intr_level level)
enum intr_level intr_enable(void)
enum intr_level intr_disable(void)


/* 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);


/* Memory */
void *memcpy(void *dest, const void * src, size_t n)
```

8. Scratch Paper (feel free to remove)

9. **Scratch Paper (feel free to remove)**