

Section 6: Synchronization and Scheduling

October 2-3, 2018

Contents

1	Warmup	2
1.1	Hello World	2
2	Vocabulary	2
3	Problems	3
3.1	Hello World Continued	3
3.2	SpaceX Problems	5
3.3	I Love You, Unconditionally!	6
3.4	Locking Up the Floopies	8
3.5	All Threads Must Die	10
3.6	Baking with Condition Variables	13

1 Warmup

1.1 Hello World

This code compiles (given a sprinkling of `#includes` etc.) but doesn't work properly. Why?

```
pthread_mutex_t lock;
pthread_cond_t cv;
int hello = 0;

void* print_hello(void* arg) {
    hello += 1;
    printf("First line (hello=%d)\n", hello);
    pthread_cond_signal(&cv);
    pthread_exit(0);
}

int main() {
    pthread_t thread;
    pthread_create(&thread, NULL, print_hello, NULL);
    while (hello < 1) {
        pthread_cond_wait(&cv, &lock);
    }
    printf("Second line (hello=%d)\n", hello);
    return 0;
}
```

This won't work because the main thread should have locked the lock before calling `pthread_cond_wait`, and the child thread should have locked the lock before calling `pthread_cond_signal`. (Also, we never initialized the lock and cv.)

2 Vocabulary

- **Lock** - Synchronization variables that provide mutual exclusion. Threads may acquire or release a lock. Only one thread may hold a lock at a time. If a thread attempts to acquire a lock that is held by some other thread, it will block at that line of code until the lock is released and it successfully acquires it. Implementations can vary.
- **Scheduler** - Routine in the kernel that picks which thread to run next given a vacant CPU and a ready queue of unblocked threads. See `next_thread_to_run()` in Pintos.
- **Priority Inversion** - If a higher priority thread is blocking on a resource (a lock, as far as you're concerned but it could be the Disk or other I/O device in practice) that a lower priority thread holds exclusive access to, the priorities are said to be inverted. The higher priority thread cannot continue until the lower priority thread releases the resource. This can be amended by implementing priority donation.
- **Priority Donation** - If a thread attempts to acquire a resource (lock) that is currently being held, it donates its effective priority to the holder of that resource. This must be done recursively until a thread holding no locks is found, even if the current thread has a lower priority than the current resource holder. (Think about what would happen if you didn't do this and a third thread with higher priority than either of the two current ones donates to the original donor.) Each thread's effective priority becomes the max of all donated priorities and its original priority.

- **Condition Variable** - A synchronization variable that provides serialization (ensuring that events occur in a certain order). A condition variable is associated with:
 - a lock (a condition variable + its lock are known together as a **monitor**)
 - some boolean condition (e.g. `hello < 1`)
 - a queue of threads waiting for the condition to be true

In order to access any CV functions **OR** to change the truthfulness of the condition, a thread must/should hold the lock. Condition variables offer the following methods:

- `cv_wait(cv, lock)` - Atomically unlocks the lock, adds the current thread to `cv`'s thread queue, and puts this thread to sleep.
- `cv_notify(cv)` - Removes one thread from `cv`'s queue, and puts it in the ready state.
- `cv_broadcast(cv)` - Removes all threads from `cv`'s queue, and puts them all in the ready state.

When a `wait()`ing thread is notified and put back in the ready state, it also re-acquires the lock before the `wait()` function returns.

When a thread runs code that may potentially make the condition true, it should acquire the lock, modify the condition however it needs to, call `notify()` or `broadcast()` on the condition's CV, so waiting threads can be notified, and finally release the lock.

Why do we need a lock anyway? Well, consider a race condition where thread 1 evaluates the condition `C` as false, then thread 2 makes condition `C` true and calls `cv.notify`, then 1 calls `cv.wait` and goes to sleep. Thread 1 might never wake up, since it went to sleep too late.

- **Hoare Semantics** - (In terms of condition variable) Wake a blocked thread when the condition is true and transfer control of the CPU and ownership of the lock to that thread immediately. This is difficult to implement in practice and generally not used despite being conceptually easier to deal with.
- **Mesa Semantics** - (In terms of condition variable) Wake a blocked thread when the condition is true, with no guarantee that the thread will execute immediately. The newly woken thread simply gets put on the ready queue and is subject to the same scheduling mechanisms as any other thread. The implication of this is that **you must check the condition with a while loop instead of an if statement because it is possible for the condition to change to false between the time the thread was unblocked and the time it takes over the CPU.**

3 Problems

3.1 Hello World Continued

Add in the necessary code to the warmup to make it work correctly.

Acquire a lock before the cv is used and release it afterwards.

```
void* print_hello(void* arg) {
    pthread_mutex_lock(&lock);
    hello += 1;
    printf("First line (hello=%d)\n", hello);
    pthread_cond_signal(&cv);
    pthread_mutex_unlock(&lock);
    pthread_exit(0);
}
```

```
}

int main() {
    pthread_t thread;
    pthread_mutex_init(&lock, 0);
    pthread_cond_init(&cv, 0);

    pthread_create(&thread, NULL, print_hello, NULL);

    pthread_mutex_lock(&lock);
    while (hello < 1) {
        pthread_cond_wait(&cv, &lock);
    }
    pthread_mutex_unlock(&lock);

    printf("Second line (hello=%d)\n", hello);
    return 0;
}
```

3.2 SpaceX Problems

Consider this program.

```
pthread_mutex_t lock = PTHREAD_MUTEX_INITIALIZER;
pthread_cond_t cv = PTHREAD_COND_INITIALIZER;

int n = 3;

void* counter(void* arg) {
    pthread_mutex_lock(&lock);
    for (n = 3; n > 0; n--)
        printf("%d\n", n);
    pthread_cond_signal(&cv);
    pthread_mutex_unlock(&lock);
}

void* announcer(void* arg) {
    while (n != 0) {
        pthread_mutex_lock(&lock);
        pthread_cond_wait(&cv, &lock);
        pthread_mutex_unlock(&lock);
    }
    printf("FALCON HEAVY TOUCH DOWN!\n");
}

int main() {
    pthread_t t1, t2;
    pthread_create(&t1, NULL, counter, NULL);
    pthread_create(&t2, NULL, announcer, NULL);
    pthread_join(t1, NULL);
    pthread_join(t2, NULL);
    return 0;
}
```

What is wrong with this code?

The lock in `announcer()` should be outside of the while loop. Or else, the announcer thread might never wake up.

3.3 I Love You, Unconditionally!

Consider the following program:

```
int subaru = 0;

int main() {
    pthread_t thread;
    pthread_create(&thread, NULL, &helper, NULL);
    pthread_yield();
    if (subaru==1) printf("I love Emilia!\n");
    else printf("I love Rem!\n");
    return 0;
}

void *helper(void* arg) {
    subaru+=1;
    pthread_exit(0);
}
```

Modify the program above such that we always print out the canonically correct line ("I love Emilia!", unfortunately). **You may only add lines.** You may not assume anything about the scheduler other than that it behaves with Mesa semantics. (In general, user programs should not depend on the scheduler and should run correctly regardless of the scheduler used.)

```
#include <stdio.h>
#include <pthread.h>

int subaru = 0;
pthread_mutex_t lock = PTHREAD_MUTEX_INITIALIZER;
pthread_cond_t cv = PTHREAD_COND_INITIALIZER;

void *helper(void* arg) {
    pthread_mutex_lock(&lock);
    subaru+=1;
    pthread_cond_signal(&cv);
    pthread_mutex_unlock(&lock);
    pthread_exit(0);
}

int main() {
    pthread_t thread;
    pthread_mutex_lock(&lock);
    pthread_create(&thread, NULL, &helper, NULL);
    pthread_yield();
    while (subaru != 1) {
        pthread_cond_wait(&cv, &lock);
    }
    if (subaru==1) {
        //SHOULD ALWAYS BE TRUE (regrettably)
        printf("I love Emilia!\n");
    } else {
        printf("I love Rem!\n");
    }
}
```

```
pthread_mutex_unlock(&lock);  
return 0;  
}
```

3.4 Locking Up the Floopies

In section 5, you may remember encountering race conditions inside of the Central Galactic Floopy Corporation's currency exchange server, which runs on top of pthreads. We said that we could make the transactions run correctly by making the balance increment/decrement atomic. The Central Galactic Floopy Corporation hires a consultant named Morty who suggests making the increment/decrement pair appear atomic by adding a lock to each account, and acquiring the locks when we run the transaction.

```
typedef struct account_t {
    pthread_mutex_t lock = PTHREAD_MUTEX_INITIALIZED;
    int balance;
    long uuid;
};

void transfer(account_t *donor, account_t *recipient, float amount) {
    // lock accounts so we can make the transfer safely
    pthread_mutex_lock(&donor->lock);
    pthread_mutex_lock(&recipient->lock);

    // check balances and make transfer if possible
    if (donor->balance < amount) {
        printf("Insufficient funds.\n");
    } else {
        donor->balance -= amount;
        recipient->balance += amount;
    }

    // unlock accounts
    pthread_mutex_unlock(&recipient->lock);
    pthread_mutex_unlock(&donor->lock);
}
```

If we use the locking code given above, will the code run correctly? Has Morty introduced a new bug into our code? Can you give an example of where this code would fail?

The locking scheme we use above will occasionally deadlock. For example, if we tried to transfer money from Bob to Alice and from Alice to Bob in separate transactions at the same time, transaction 1 will acquire Bob's lock and wait on Alice's lock, while transaction 2 acquires Alice's lock while waiting on Bob's lock.

Can you modify the code above to resolve this bug?

```
typedef struct account_t {
    pthread_mutex_t lock = PTHREAD_MUTEX_INITIALIZED;
    int balance;
    long uuid;
};

void transfer(account_t *donor, account_t *recipient, float amount) {
    // lock accounts so we can make the transfer safely

    if (donor->uuid <= recipient->uuid) {

        pthread_mutex_lock(&donor->lock);
        pthread_mutex_lock(&recipient->lock);
    } else {
        pthread_mutex_lock(&recipient->lock);
        pthread_mutex_lock(&donor->lock);
    }

    // check balances and make transfer if possible
    if (donor->balance < amount) {
        printf("Insufficient funds.\n");
    } else {
        donor->balance -= amount;
        recipient->balance += amount;
    }

    // unlock accounts
    pthread_mutex_unlock(&recipient->lock);
    pthread_mutex_unlock(&donor->lock);
}
```

3.5 All Threads Must Die

You have three Pintos threads with the associated priorities shown below. They each run the functions with their respective names.

Tyrion : 4
Ned: 5
Gandalf: 11

Assume upon execution that all threads are unblocked and begin at the top of their code blocks. The operating system runs with a preemptive priority scheduler. You may assume that `set_priority` commands are atomic. (Note: The following uses references to Pintos locks and data structures.)

```
struct list brace_yourself; // pintos list. Assume it's already initialized and populated.
struct lock midterm;       // pintos lock. Already initialized.
struct lock is_coming;
```

```
void tyrion(){
    thread_set_priority(12);
    lock_acquire(&midterm);
    lock_release(&midterm);
    thread_exit();
}

void ned(){
    lock_acquire(&midterm);
    lock_acquire(&is_coming);
    list_remove(list_head(brace_yourself));
    lock_release(&midterm);
    lock_release(&is_coming);
    thread_exit();
}

void gandalf(){
    lock_acquire(&is_coming);
    thread_set_priority(3);
    while (thread_get_priority() < 11) {
        printf("YOU .. SHALL NOT .. PAAASS!!!!!!");
        timer_sleep(20);
    }
    lock_release(&is_coming);
    thread_exit();
}
```

What is the output of this program when there is no priority donation? Trace the program execution and number the lines in the order in which they are executed.

```
void tyrion(){
5   thread_set_priority(12);
6   lock_acquire(&midterm); //blocks
    lock_release(&midterm);
    thread_exit();
```

```

}

void ned(){
3   lock_acquire(&midterm);
4   lock_acquire(&is_coming); //blocks
   list_remove(list_head(brace_yourself));
   lock_release(&midterm);
   lock_release(&is_coming);
   thread_exit();
}

void gandalf(){
1   lock_acquire(&is_coming);
2   thread_set_priority(3);
7   while (thread_get_priority() < 11) {
8       printf("YOU .. SHALL NOT .. PAAASS!!!!!!"); //repeat till infinity
9       timer_sleep(20);
   }
   lock_release(&is_coming);
   thread_exit();
}

```

Gandalf, as you might expect, endlessly prints "YOU SHALL NOT PASS!!" every 20 clock ticks or so.\\

What is the output and order of line execution if priority donation was implemented? Draw a diagram of the three threads and two locks that shows how you would use data structures and struct members (variables and pointers, etc) to implement priority donation for this example.

```

void tyrion(){
8   thread_set_priority(12);
9   lock_acquire(&midterm); //blocks
   lock_release(&midterm);
   thread_exit();
}

void ned(){
3   lock_acquire(&midterm);
4   lock_acquire(&is_coming); //blocks
11  list_remove(list_head(brace_yourself)); //KERNEL PANIC
   lock_release(&midterm);
   lock_release(&is_coming);
   thread_exit();
}

void gandalf(){
1   lock_acquire(&is_coming);
2   thread_set_priority(3);
5   while (thread_get_priority() < 11) { //priority is 5 first, but 12 at some later loop
6       printf("YOU .. SHALL NOT .. PAAASS!!!!!!");
7       timer_sleep(20);
   }
10  lock_release(&is_coming);
}

```

```
    thread_exit();  
}
```

It turns out that Gandalf generally does mean well. Donations will make Gandalf allow you to pass. At some point Gandalf will sleep on a timer and leave Tyrion alone in the ready queue. Tyrion will run even though he has a lower priority (Gandalf has a 5 donated to him) Tyrion then sets his priority to 12 and chain-donates to Gandalf. Gandalf breaks his loop. Ned unblocks after Gandalf releases the `is_coming` lock. However, allowing Ned to remove the head of a list will trigger an ASSERT failure in `lib/kernel/list.c`.

Gandalf will print `YOU SHALL NOT PASS` at least once. Then Ned will get beheaded and cause a kernel panic that crashes Pintos.

3.6 Baking with Condition Variables

A number of people are trying to bake cakes. Unfortunately, they each know only one skill, so they need to all work together to bake cakes. Use independent threads (one person is one thread) which communicate through condition variables to solve the problem. A skeleton has been provided, fill in the blanks to make the implementation work.

A cake requires:

- 1 cake batter
- 2 eggs

Instructions:

1. Add ingredients to bowl
2. Heat bowl (it's oven-safe)
3. Eat the cake, clean out the bowl, and go back to step 1

Requirements:

- Don't start heating the cake in the oven unless there are exactly the right number of ingredients in the bowl.
- Don't add raw ingredients to a currently-baking cake or a finished cake.
- Don't eat the cake unless it's done baking.
- Given enough time, the code should bake an unbounded number of cakes, and should never stop.

(Hint: start with the cake eater; we've filled out a bit more of his duties.)

```
int numBatterInBowl = 0;
int numEggInBowl = 0;
bool readyToEat = false;
pthread_mutex_t lock;
pthread_cond_t needIngredients;
pthread_cond_t readyToBake;
pthread_cond_t startEating;

void* batterAdder(void* arg)
{
    pthread_mutex_lock(&lock);
    while (1) {
        while (numBatterInBowl || readyToEat) {

            -----
        }
        addBatter(); // Sets numBatterInBowl += 1
        pthread_cond_signal(&readyToBake);
    }
}
```

```

void* eggBreaker(void* arg)
{
    pthread_mutex_lock(&lock);
    while (1) {

        ----- {

            -----
        }
        addEgg(); // Sets numEggInBowl += 1

        -----
    }
}

void* bowlHeater(void* arg)
{
    pthread_mutex_lock(&lock);
    while (1) {

        ----- {

            -----
        }
        heatBowl(); // Sets readyToEat = true, numBatterInBowl = 0, numEggInBowl = 0

        -----
    }
}

void* cakeEater(void* arg)
{
    pthread_mutex_lock(&lock);
    while (1) {

        ----- {

            -----
        }
        eatCake(); // Sets readyToEat = false and cleans the bowl for another cake

        -----
    }
}

int main(int argc, char *argv[])
{
    // Initialize mutex and condition variables
    // Start threads: 1 batterAdder, 2 eggBreakers, 1 bowlHeater, and 1 cakeEater
    // main() sleeps forever
}

```

```
// P3_5.c
#include <stdio.h>
#include <pthread.h>
#include "cake.h"
#include <unistd.h>

#define TRUE 1
#define FALSE 0

int numBatterInBowl = 0;
int numEggInBowl = 0;
int readyToEat = FALSE;
pthread_mutex_t lock;
pthread_cond_t needIngredients;
pthread_cond_t readyToBake;
pthread_cond_t startEating;

void addBatter() {
    numBatterInBowl += 1;
    printf ("Added batter\n");
}

void addEgg() {
    numEggInBowl += 1;
    printf ("Added egg\n");
}

void heatBowl() {
    readyToEat = TRUE;
    numBatterInBowl = 0;
    numEggInBowl = 0;
    printf ("Heated bowl\n");
}

void eatCake() {
    readyToEat = FALSE;
    printf ("Anyway, this cake is great.\n"
           "It's so delicious and moist.\n"
           "%s", CAKE);
}

void* batterAdder(void* arg)
{
    pthread_mutex_lock(&lock);
```

```

while (1) {
    while (numBatterInBowl || readyToEat) {
        pthread_cond_wait(&needIngredients, &lock);
    }
    addBatter(); // Sets numBatterInBowl += 1
    pthread_cond_signal(&readyToBake);
}
}

void* eggBreaker(void* arg)
{
    pthread_mutex_lock(&lock);
    while (1) {
        while (numEggInBowl >= 2 || readyToEat) {
            pthread_cond_wait(&needIngredients, &lock);
        }
        addEgg(); // Sets numEggInBowl += 1
        pthread_cond_signal(&readyToBake);
    }
}

void* bowlHeater(void* arg)
{
    pthread_mutex_lock(&lock);
    while (1) {
        while (numBatterInBowl != 1 || numEggInBowl != 2) {
            pthread_cond_wait(&readyToBake, &lock);
        }
        heatBowl(); // Sets readyToEat = true, numBatterInBowl = 0, numEggInBowl = 0
        pthread_cond_signal(&startEating);
    }
}

void* cakeEater(void* arg)
{
    pthread_mutex_lock(&lock);
    while (1) {
        while (!readyToEat) {
            pthread_cond_wait(&startEating, &lock);
        }
        eatCake(); // Sets readyToEat = false and cleans the bowl for another cake
        pthread_cond_broadcast(&needIngredients);
    }
}

int main(int argc, char *argv[])
{
    pthread_mutex_init(&lock, 0);
    pthread_cond_init(&needIngredients, 0);
}

```



```

pthread_cond_init(&readyToBake, 0);
pthread_cond_init(&startEating, 0);

pthread_t thread1, thread2, thread3, thread4, thread5;
pthread_create(&thread1, NULL, batterAdder, NULL);
pthread_create(&thread2, NULL, eggBreaker, NULL);
pthread_create(&thread3, NULL, eggBreaker, NULL);
pthread_create(&thread4, NULL, bowlHeater, NULL);
pthread_create(&thread5, NULL, cakeEater, NULL);

while (1) {
    sleep (1000);
}

return 0;
}

// cake.h
#define CAKE "
                ,: /+/-
                \n" \
"
                /M/
                ., -=; //;- \n" \
"
                .:/= ;MH/,
                ,= /+%$XH@MM#@: \n" \
"
                -$##@+$###@H@MMM#####H:.
                -/H#\n" \
"
                .,H@H@ X#####@ -H#####@+-
                -+H#####@X\n" \
"
                .,@##H;
                +XM##M/,
                =%@###@X;- \n" \
"X%-
                :M#####$.
                .:%M###@%: \n" \
"M##H,
                +H@@@$/-.
                ,;$M###@%,
                -\n" \
"M###M=,
                ,---,
                .-%H###M$:
                ,+@##\n" \
"@#####@/
                :%H##@$- \n" \
"M#####H,
                ;HM##M$= \n" \
"#####.
                .=$M##M$= \n" \
"#####H..;XM##M$=
                .:+\n" \
"M#####@%=
                =+@MH%\n" \
"@#####M/.
                =+H#X%= \n" \
"=+M#####M,
                ,/X#H+:,
                \n" \
"
                .;XM#####H=
                ,/X#H+:;
                \n" \
"
                .+=HM#####M+/+HM@+=.
                \n" \
"
                ,:/%XM###H/.
                \n" \
"
                ,.:=-.
                \n"

```