



CS/COE 1550 – Introduction to Operating Systems

Project 2: Process Synchronization

Due: Tuesday, March 3, 2020 @11:59pm

Late: Thursday, March 5, 2020 @11:59pm with 10% reduction per late day

Table of Contents

PROJECT OVERVIEW	2
PROJECT DETAILS	2
VISITORARRIVES() AND TOURGUIDEARRIVES()	2
TOURMUSEUM() AND OPENMUSEUM()	3
VISITORLEAVES() AND TOURGUIDELEAVES()	3
VISITOR ARRIVAL PROCESS.....	4
TOUR GUIDE ARRIVAL PROCESS	4
REQUIREMENTS	4
TESTING	4
PROGRAM AND OUTPUT SPECS	4
SHARED MEMORY IN OUR SIMULATION	6
BUILDING AND RUNNING MUSEUMSIM	6
DEBUGGING.....	6
SUBMISSION	7
GRADING SHEET/RUBRICS	7



CS/COE 1550 – Introduction to Operating Systems

Project Overview

Anytime we share data between two or more processes or threads, we run the risk of having race conditions and deadlocks. In race conditions, the shared data could become corrupted. In order to avoid these situations, we have discussed various mechanisms to ensure that one process's critical regions are guarded from another's.

One place that we might use parallelism is to simulate real-world situations that involve multiple independently acting entities, such as people. In this project, you will use the semaphore implementation that you finished in Project 1¹ to model the **safe museum tour problem**, whereby visitors and museum tour guides synchronize so that:

- a visitor cannot tour a museum without a tour guide,
- a tour guide cannot open the museum without a visitor,
- a tour guide leaves when no more visitors are in the museum,
- a tour guide cannot leave until all visitors in the museum leave,
- at most **two** tour guides can be in the museum at a time, and
- each tour guide provides a tour for **at most ten** visitors.

Your job is to write a program that (a) always satisfies the above constraints and (b) under no conditions will cause a deadlock to occur. A deadlock happens, for example, when the museum is empty, a tour guide and a visitor arrive, and they stay waiting outside the museum forever.

Project Details

You are to write (a) the visitor process, (b) the tour guide process, and (c) **six** functions called by these processes: (c1) `visitorArrives()`, (c2) `tourMuseum()`, (c3) `visitorLeaves()`, (c4) `tourguideArrives()`, (c5) `openMuseum()`, and (c6) `tourguideLeaves()`. You will also write two processes, (d) one for simulating tour guides' arrival and (e) the other for simulating visitors' arrival.

`visitorArrives()` and `tourguideArrives()`

In order to open a museum for tours, there need to be at least one tour guide and one visitor. `visitorArrives()` is called by visitor processes and `tourguideArrives()` is called by tour guide processes. Both functions must **block** until a tour guide **and** a visitor both arrive.

An arriving visitor must wait if

- the museum is closed or

¹ If you haven't finished Project 1 or are not sure if your solution is correct, please download and use the modified kernel from CourseWeb.



CS/COE 1550 – Introduction to Operating Systems

- the maximum number of visitors has been reached

An arriving tour guide must wait if

- the museum is closed, and no visitor is waiting outside or
- the museum is open, and two tour guides are inside the museum.

While only one tour guide is inside the museum,

- up to ten visitors may arrive (i.e., call `visitorArrives()`) and
- a second tour guide may open the museum (i.e., call `openMuseum()`).

When a tour guide arrives, the following message is printed to the screen:

Tour guide %d arrives at time %d.

When a visitor arrives, the following message is printed to the screen:

Visitor %d arrives at time %d.

`tourMuseum()` and `openMuseum()`

After `tourguideArrives()` returns, the tour guide immediately calls `openMuseum()`. After `visitorArrives()` returns, the visitor immediately calls `tourMuseum()`. Each visitor takes **2 seconds** to tour the museum (you can implement that by calling `nanosleep` or `sleep`). A visitor inside `tourMuseum()` must not block another visitor who is also inside `tourMuseum()`.

When a tour guide opens the museum, the following message is printed to the screen:

Tour guide %d opens the museum for tours at time %d.

When a visitor tours the museum, the following message is printed to the screen:

Visitor %d tours the museum at time %d.

`visitorLeaves()` and `tourguideLeaves()`

Tour guides that are inside the museum cannot leave until all visitors inside the museum leave.

When a tour guide leaves, the following message should be printed to the screen:

Tour guide %d leaves the museum at time %d

When a visitor leaves, the following message should be printed to the screen:

Visitor %d leaves the museum at time %d



CS/COE 1550 – Introduction to Operating Systems

Visitor Arrival Process

The visitor arrival process creates m visitor processes. The number of visitors, m , is read from a command-line argument (e.g., `./museumsim -m 10`). Visitors arrive in bursts. When a visitor arrives, there is a pv (e.g., 70%) chance that another visitor is immediately arriving after her, but once no visitors arrive, there is a dv seconds (e.g., 20 seconds) delay before any new visitor arrives. The first visitor always arrives at time 0. The probability pv and the delay dv are to be read from the command-line (e.g., `./museumsim -pv 70 -dv 20`).

Tour guide Arrival Process

The tour guide arrival process creates k tour guide processes. The number of tour guides, k , is read from a command-line argument (e.g., `./museumsim -k 10`). Tour guides arrive in bursts. When a tour guide arrives, there is a pg (e.g., 30%) chance that another tour guide is immediately arriving after² her, but once no tour guides arrive, there is a dg second (e.g., 30 seconds) delay before any new tour guide arrives. The first tour guide always arrives at time 0. The probability pa and the delay da are to be read from the command-line (e.g., `./museumsim -pa 30 -da 30`).

Requirements

Your solution must use semaphores, should not use busy waiting, and should be deadlock-free.

Testing

Make sure to run various test cases against your solution; for instance, create k tour guides and m visitors (with various values of k and m , for example $k > m$, $m > k$, $k == m$), different values for the probabilities and delays, etc. Note that the exact order of the output messages can vary, within certain boundaries. These boundaries are checked by the autograder scripts.

Program and Output Specs

Create a program, `museumsim`, which runs the simulation. Your program should run as follows.

- Fork a process for visitor arrival and a process for tour guide arrival, each in turn forking visitor processes and tour guide processes, respectively, at the appropriate times as specified by the command-line arguments.
- To get an 80% chance of something, you can generate a random number modulo 100 and see if its value is less than 80. It's like flipping an unfair coin. You may refer to CS 449 materials for how to generate a random number.

² Recall that at most two tour guides can be in the museum at a time.



CS/COE 1550 – Introduction to Operating Systems

- Have the following command-line arguments:
 - -m: number of visitors
 - -k: number of tour guides
 - -pv: probability of a visitor immediately following another visitor
 - -dv: delay in seconds when a visitor does not immediately follow another visitor
 - -sv: random seed for the visitor arrival process
 - -pg: probability of a tour guide immediately following another tour guide
 - -dg: delay in seconds when a tour guide does not immediately follow another tour guide
 - -sg: random seed for the tour guide arrival process
- Make sure that your **output** shows all the necessary events. You should sequentially number each visitor and tour guide starting from 0. Visitor numbers are independent of tour guide numbers. When the museum is empty, your program should display:

The museum is now empty.

- Print out messages in the form:

The museum is now empty.

Tour guide %d arrives at time %d.

Visitor %d arrives at time %d.

Tour guide %d opens the museum for tours at time %d.

Visitor %d tours the museum at time %d.

Visitor %d leaves the museum at time %d.

Tour guide %d leaves the museum at time %d.

- The printed time is in **seconds since the start of the program**. You should use the function `gettimeofday` and use both the seconds and microseconds fields in `struct timeval` in your calculation of the number of elapsed seconds.



CS/COE 1550 – Introduction to Operating Systems

Shared Memory in our simulation

To make our shared data and our semaphores, what we need is for multiple processes to be able to share the same memory region. We can ask for N bytes of RAM from the OS directly by using `mmap()`:

```
void *ptr = mmap(NULL, N, PROT_READ|PROT_WRITE, MAP_SHARED|MAP_ANONYMOUS, 0, 0);
```

The return value will be the address of the start of the allocated memory. We can then steal portions of the allocated memory to hold our variables much like the `malloc()` project from CS 0449. For example, if we wanted two integers to be stored in the allocated region, we could do the following to allocate and initialize them.

```
int *first_sem;
int *second_sem;
first = (int *) ptr;
second_sem = first_sem + 1;
*first_sem = 0;
*second_sem = 0;
```

At this point we have one process and some RAM that contains our variables. But we now need to share that memory region/variables with other processes. The good news is that a `mmap`'ed region (with the `MAP_SHARED` flag) remains accessible in the child process after a `fork()`. Therefore, do the `mmap()` in `main` before `fork()` and then use the variables in the appropriate way afterwards.

Building and Running `museumsim`

Please use the instructions in Project 1 for building and running the test program (named `trafficsim` in Project 1) to build and run `museumsim`. Make sure that the QEMU VM is running the modified kernel (with the `down` and `up` syscalls). Whenever you make a change to `museumsim.c`, you only to `scp` it into QEMU but you don't need to reboot QEMU.

Debugging

Please use the following steps to run `gdb` inside the QEMU virtual machine:

Inside the QEMU VM:

```
scp PITT_ID@thoth.cs.pitt.edu:/u/OSLab/original/gdb/* .
mv gdb /bin/
mv libncurses.so.5 /lib
mv libtinfo.so.5 /lib

gdb <program name>
```

On `thoth`:



CS/COE 1550 – Introduction to Operating Systems

You will have to add the -g flag to the compilation command (typed on thoht) for the museumsim program.

To run valgrind inside the QEMU VM:

On QEMU:

```
scp -r PITT_ID@thoth.cs.pitt.edu:/u/OSLab/original/valgrind/* .
mv bin/* /bin/
mv lib/* /lib/
echo export VALGRIND_LIB=/lib/valgrind >> ~/.bashrc
bash
valgrind ./<program name with command-line arguments>
```

Submission

You need to submit the following to Gradescope by the deadline:

- Your well-commented museumsim.c program's source,
- the file sem.h from Project 1 or from the provided kernel, and
- a brief, intuitive explanation of why (or why not) your solution is fair (maximum 10 visitors per tour guide), as well as deadlock- and starvation-free.

Grading Sheet/Rubrics

Item	Grade
Test cases on the autograder	80%
Comments and style	10%
Explanation report	10%

Please note that the score that you get from the autograder is not your final score. We still do manual grading. We may discover bugs and mistakes that were not caught by the test scripts and take penalty points off. Please use the autograder only as a tool to get immediate feedback and discover bugs in your program. Please note that certain bugs (e.g, deadlocks) in your program may or may not manifest when the test cases are run on your program. It may be the case that the same exact code fails in some tests then the same code passes the same tests when resubmitted. The fact that a test once fails should make you try to debug the issue not to resubmit and hope that the situation that caused the bug won't happen the next time around.