[2024-Fall] Operating System
Project #2 Instruction
Implementation and analysis of the virtual CPU and its scheduling

Lecturer	Gyeongsik Yang (g_yang@korea.ac.kr) System Software Lab. **https://ss.korea.ac.kr**
TAs	Jinwoo Jeong, Jongin Choi (osta@os.korea.ac.kr)

Topic	Implement custom virtual CPU object in kernel and analyze process scheduling policies on it
Due date(submission)	**2024. 11. 12. (Tuesday) PM 11:59**
Recommended environment	Ubuntu 18.04.02 (64 bit), Linux kernel 4.20.11 (VM environment)

Contents

1. Purpose

This project aims to create a virtual CPU object in kernel space and analyze the process scheduling policies applied to it. We will implement a virtual CPU object that functions similarly to real-world CPUs. The virtual CPU object will be created and exist in kernel space. This object can be used (occupied) by user applications through a system call. From the perspective of the virtual CPU object, it should be initialized, wait for process requests for CPU usage, be allocated (used) by processes as requested under various scheduling policies, and maintain its state within the kernel.

2. Implementation environment

We will call the new virtual CPU object "**ku_cpu**." Note that in our system, we assume that our system contains only one **ku_cpu** (CPU core).

i) User space

A. Implement a user process that requests the use of a virtual CPU.
B. The processes use a system call to request the virtual CPU.
C. Print both response time and wait time

ii) Kernel space

A. Implement a virtual CPU object that behaves like a real CPU.
B. It can be implemented similarly to **Project #1** using a system call handler.
C. Scheduling polices
   - FCFS
   - SJF

* See details for each below.

3. User space: user-space program for requesting ku_cpu

Our user program takes several arguments for its execution:

1. the execution duration (in seconds)
2. the starting delay (in seconds)
3. the name of the process

For example, we will run the program by giving the arguments "7 1 A" meaning that the process of name "A" will request the use of **ku_cpu** after 1 second from its execution and will use **ku_cpu** 7 seconds.

We assume that when the request on the **ku_cpu** is succeeded by system call, 0.1 seconds of the execution duration is well handled. After using **ku_cpu** for the given execution duration, the process should print the total waiting time (with round-up).

Figure 1 shows the user space program implementation. Initially, **ku_cpu** system call number is defined with the **#define** statement. When calling the system call, the user program should pass 1) the process name (name) and 2) the remaining job time (job) as arguments. The user process uses the return value to determine whether the request has been accepted or not. See the comments in the source code for more details.

kucpu_run.c 파일임

```c
#include <unistd.h>
#include <stdio.h>
#include <string.h>
#include <stdlib.h>

#define KU_CPU 339 // define syscall number

int main(int argc, char ** argv){
    int jobTime;
    int delayTime;
    char name[4];
    int wait = 0;

    if (argc < 4){
        printf("\nInsufficient Arguments..\n");
        return 1;
    }

    /*  first argument: job time (second)
        second argument: delay time before execution (second)
        third argument: process name
    */

   jobTime = atoi(argv[1]);
   delayTime = atoi(argv[2]);
   strcpy(name,argv[3]);

   // wait for 'delayTime' seconds before execution
   sleep(delayTime);
   printf("\nProcess %s : I will use CPU by %ds.\n",name,jobTime);
   jobTime *= 10; // exectute system call in every 0.1 second

   // continue requesting the system call as long as the jobTime remains
   while(jobTime){
    // if request is rejected, increase wait time
    if (!syscall(KU_CPU, name, jobTime)) jobTime--;
    else wait ++;
    usleep(100000); // delay 0.1 second
   }

   syscall(KU_CPU,name,0);
   printf("\nProcess %s : Finish! My total wait time is %ds. ",name, (wait+5)/10);
   return 0;
}
```

Figure 1: user-space program implementation without response time

Note that the provided source code is incomplete. You should modify this code to **print both response time and wait time**.

4. ku_cpu: virtual CPU object

4.1 Operation sequence example

The following is the high-level operation sequence of **ku_cpu with FCFS:**

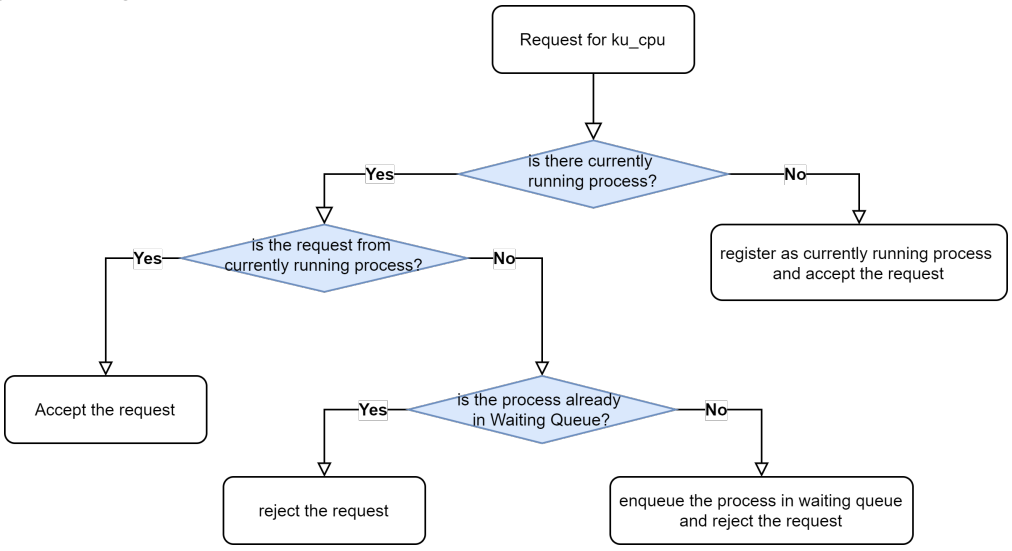

Figure 2: flow of **ku_cpu** usage.

1. The **ku_cpu** is implemented in kernel-space. User processes use a system call to request the use of the **ku_cpu** by providing two parameters: the process name, the time duration for occupying the **ku_cpu**. Upon this request, the **ku_cpu** in the system call handler executes the following tasks.
2. The **ku_cpu** keeps track of which process is currently running (occupying the **ku_cpu**):
   A. If the **ku_cpu** is idle (not occupied by any process), the first process that requests the **ku_cpu** is allowed to occupy it. In this case, the **ku_cpu** internally sets this process as the running process. Then, it returns 0 to notify the user process of successful **ku_cpu** occupation.
   B. If the process currently requesting the **ku_cpu** is the same one that is already occupying it, the system will return 0 to confirm the successful continuation of **ku_cpu** usage.
   C. If a different process requests the **ku_cpu** while it is occupied, the different process is added to the "waiting queue." However, if this process is already in the waiting queue, it will not be enqueued again. The system then returns 1 to notify the user process of the waiting status, indicating a temporary inability to occupy the **ku_cpu** immediately.
3. If the process currently using the **ku_cpu** has no more tasks remaining (the allocated duration has expired), it implies that the process has finished its execution on the **ku_cpu**. Subsequently, the next process is dequeued (popped) from the waiting queue and assigned the **ku_cpu**. If there is no process in the waiting queue, the **ku_cpu** reverts to an idle state.

4.2 Data structure (or variables) for ku_cpu

To implement the **ku_cpu**, the following structures and variables are essential. You may use any type of implementation or data structures you required:

1. Waiting queue (processes waiting to use **ku_cpu**)
2. Variable to keep track of the current process occupying the **ku_cpu**
3. PID of the process being processed by **ku_cpu**

   To implement the **ku_cpu**, we should distinguish and identify the processes that called the system call. In Linux, a process is created through the **fork()** system call, and kernel assigns a unique PID value to each process. This PID is usually an integer value of 3 to 5 digits. The type of PID is **pid_t**, but it can be handled like an integer, and it can also be printed using **printk**.

   We can obtain the PID from the PCB (process control block) structure in Linux, which is **task_struct**. Kernel provides a pointer to the **task_struct** structure of the current context. By adding the file **linux/sched.h** as a header file, we can get the PID value.

```c
#include <linux/sched.h>
#include <linux/kernel.h>
#include <linux/syscalls.h>
#include <linux/linkage.h>
#include <linux/slab.h>

SYSCALL_DEFINE1(os2024_pid_print, char*, name) {
        pid_t pid = current->pid;

        printk("Process name: %s pid: %d\n",name,pid);

        return 0;

}
```

Figure 3: print pid value

4.3 Example implementation of ku_cpu

```c
SYSCALL_DEFINE2(os2024_ku_cpu, char*, name, int, jobTime){
        // store pid of current process as pid_t type
        job_t newJob = {current->pid, jobTime};

        // register the process if virtual CPU is idle
        if (now==IDLE) now = newJob.pid;

        // If the process that sent the request is currently using virtual CPU
        if (now == newJob.pid) {
                // If the job has finished
                if (jobTime == 0) {
                        printk("Process Finished: %s\n",name);
                        // if queue is empty, virtual CPU becomes idle
                        if (ku_is_empty()) now = IDLE;
                        // if not, get next process from queue
                        else now = ku_pop();
                }
                else printk("Working: %s\n", name);
                // request accepted
                return 0;
        }
        else {
                // if the requset is not from currently handling process
                if (ku_is_new_id(newJob.pid)) ku_push(newJob);
                printk("Working Denied:%s \n", name);
        }
        // request rejected
        return 1;
}
```

Figure 4: **ku_cpu** (by system call handler) implementation example.

Figure 4 is an example source code of a **ku_cpu** system call. Please note that the provided source code is incomplete and will not function if run as-is. It lacks several critical components, such as external or global variables and the implementation of the waiting queue. It is your responsibility to complete the remaining parts of the implementation:

4.4 Scheduling policies to implement

The above explanation is based on FCFS (first-come first-served) policy, known for its simplicity and ease of implementation. However, it has disadvantages as well. In this project, you are to start by fully implementing the FCFS scheduling policy. Once this is accomplished, you are expected to implement the additional scheduling policy:

1. First-come first-served (FCFS)
2. Shortest job first (SJF)

* For detailed explanations and characteristics of each scheduling policy, please refer to our lecture notes.

5. Experiment and analysis

5.1 Experiment configuration

To validate our implementation, we will conduct experiments and report the results. For the experiment, run the user program as three separate processes concurrently and simultaneously. Assume that we have compiled the user program and obtained an executable file named **p**.

Then, create a new file named "**run**" with the following commands. You can adjust the execution duration or start delay to effectively reveal the differences between the scheduling policies.

- For both FCFS and SJF

```
./p 7 0 A &
./p 5 1 B &
./p 3 2 C &
```

The "&" symbol at the end of each command instructs the system to execute the process in the background in Linux. This allows us to execute all three **p** programs concurrently. You should try out these commands in the terminal to ensure you understand the procedure.

Next, enter the following command in the terminal to change the execution permissions of the **run** file. This command allows all users in Linux to read, write, and execute the **run** file:

```
$ chmod 777 run
```

When executing the **run with FCFS**, process A will start immediately with a duration of 7 seconds. One second later, process B will start and run for 5 seconds. Following another second, process C will start and run for 3 seconds.

* You can modify this **run** file to conduct your own test for the report.

5.2 Expected results

Your logs do not have to match the expected results exactly, but they should contain the necessary scheduling information, such as the wait time, response time for each user process, and whether the **ku_cpu** granted or denied usage.

- FCFS (user process log from console)

```
osta@osta-VirtualBox:~/project/n2$ ./run
osta@osta-VirtualBox:~/project/n2$
Process A : I will use CPU by 7s.

Process B : I will use CPU by 5s.

Process C : I will use CPU by 3s.

Process A : Finish! My response time is 0s and My total wait time is 0s.
Process B : Finish! My response time is 6s and My total wait time is 6s.
Process C : Finish! My response time is 10s and My total wait time is 10s.
```

- FCFS (kernel log from **dmesg**)

```
[목  7월 18 16:58:02 2024] Working: A
[목  7월 18 16:58:02 2024] Working: A
[목  7월 18 16:58:03 2024] Working: A
[목  7월 18 16:58:03 2024] Working: A
[목  7월 18 16:58:03 2024] Working: A
[목  7월 18 16:58:03 2024] Working: A
[목  7월 18 16:58:03 2024] Working: A
[목  7월 18 16:58:03 2024] Working: A
[목  7월 18 16:58:03 2024] Working Denied:B
[목  7월 18 16:58:03 2024] Working: A
[목  7월 18 16:58:03 2024] Working Denied:B
[목  7월 18 16:58:03 2024] Working: A
[목  7월 18 16:58:03 2024] Working Denied:B
[목  7월 18 16:58:03 2024] Working: A
[목  7월 18 16:58:03 2024] Working Denied:B
[목  7월 18 16:58:03 2024] Working: A
[목  7월 18 16:58:04 2024] Working Denied:B
[목  7월 18 16:58:04 2024] Working: A
[목  7월 18 16:58:04 2024] Working Denied:B
[목  7월 18 16:58:04 2024] Working: A
[목  7월 18 16:58:04 2024] Working Denied:B
[목  7월 18 16:58:04 2024] Working: A
[목  7월 18 16:58:04 2024] Working Denied:B
[목  7월 18 16:58:04 2024] Working: A
```

```
[목  7월 18 16:58:09 2024] Working: A
[목  7월 18 16:58:09 2024] Working Denied:B
[목  7월 18 16:58:09 2024] Working Denied:C
[목  7월 18 16:58:09 2024] Working: A
[목  7월 18 16:58:09 2024] Working Denied:B
[목  7월 18 16:58:09 2024] Working Denied:C
[목  7월 18 16:58:09 2024] Process Finished: A
[목  7월 18 16:58:09 2024] Working: B
[목  7월 18 16:58:09 2024] Working Denied:C
[목  7월 18 16:58:09 2024] Working: B
[목  7월 18 16:58:09 2024] Working Denied:C
[목  7월 18 16:58:10 2024] Working: B
[목  7월 18 16:58:10 2024] Working Denied:C
[목  7월 18 16:58:10 2024] Working: B
[목  7월 18 16:58:10 2024] Working Denied:C
[목  7월 18 16:58:10 2024] Working: B
[목  7월 18 16:58:10 2024] Working Denied:C
[목  7월 18 16:58:10 2024] Working: B
[목  7월 18 16:58:10 2024] Working Denied:C
[목  7월 18 16:58:10 2024] Working: B
[목  7월 18 16:58:10 2024] Working Denied:C
[목  7월 18 16:58:10 2024] Working: B
[목  7월 18 16:58:10 2024] Working Denied:C
[목  7월 18 16:58:10 2024] Working: B
```

- SJF (user process log from console, kernel log omitted)

```
osta@osta-VirtualBox:~/project/n2$ ./sjf_run
osta@osta-VirtualBox:~/project/n2$
Process A : I will use CPU by 7s.

Process B : I will use CPU by 5s.

Process C : I will use CPU by 3s.

Process A : Finish! My response time is 0s and My total wait time is 0s.
Process C : Finish! My response time is 5s and My total wait time is 5s.
Process B : Finish! My response time is 9s and My total wait time is 9s.
```

- SJF (kernel log from **dmesg**)

```
[목  7월 18 17:28:09 2024] Working: A
[목  7월 18 17:28:09 2024] Working: A
[목  7월 18 17:28:09 2024] Working: A
[목  7월 18 17:28:09 2024] Working: A
[목  7월 18 17:28:09 2024] Working: A
[목  7월 18 17:28:09 2024] Working: A
[목  7월 18 17:28:09 2024] Working: A
[목  7월 18 17:28:09 2024] Working: A
[목  7월 18 17:28:09 2024] Working: A
[목  7월 18 17:28:09 2024] Working Denied: B
[목  7월 18 17:28:09 2024] Working: A
[목  7월 18 17:28:10 2024] Working Denied: B
[목  7월 18 17:28:10 2024] Working: A
[목  7월 18 17:28:10 2024] Working Denied: B
[목  7월 18 17:28:10 2024] Working: A
[목  7월 18 17:28:10 2024] Working Denied: B
[목  7월 18 17:28:10 2024] Working: A
[목  7월 18 17:28:10 2024] Working Denied: B
[목  7월 18 17:28:10 2024] Working: A
[목  7월 18 17:28:10 2024] Working Denied: B
[목  7월 18 17:28:10 2024] Working: A
[목  7월 18 17:28:10 2024] Working Denied: B
[목  7월 18 17:28:10 2024] Working: A
[목  7월 18 17:28:10 2024] Working Denied: B
[목  7월 18 17:28:10 2024] Working: A
[목  7월 18 17:28:10 2024] Working Denied: B
[목  7월 18 17:28:10 2024] Working: A
[목  7월 18 17:28:10 2024] Working Denied: B
[목  7월 18 17:28:10 2024] Working: A
[목  7월 18 17:28:10 2024] Working Denied: C
[목  7월 18 17:28:10 2024] Working Denied: B
[목  7월 18 17:28:11 2024] Working: A
[목  7월 18 17:28:11 2024] Working Denied: C
[목  7월 18 17:28:11 2024] Working Denied: B
[목  7월 18 17:28:11 2024] Working: A
[목  7월 18 17:28:11 2024] Working Denied: C
[목  7월 18 17:28:11 2024] Working Denied: B
[목  7월 18 17:28:11 2024] Working: A
[목  7월 18 17:28:11 2024] Working Denied: C
[목  7월 18 17:28:11 2024] Working Denied: B
[목  7월 18 17:28:11 2024] Working: A
[목  7월 18 17:28:11 2024] Working Denied: C
[목  7월 18 17:28:11 2024] Working Denied: B
[목  7월 18 17:28:11 2024] Working: A
[목  7월 18 17:28:11 2024] Working Denied: C
```

```
[목  7월 18 17:28:18 2024] Working Denied: B
[목  7월 18 17:28:18 2024] Working: C
[목  7월 18 17:28:18 2024] Working Denied: B
[목  7월 18 17:28:19 2024] Working: C
[목  7월 18 17:28:19 2024] Working Denied: B
[목  7월 18 17:28:19 2024] Working: C
[목  7월 18 17:28:19 2024] Working Denied: B
[목  7월 18 17:28:19 2024] Process Finished: C
[목  7월 18 17:28:19 2024] Working: B
[목  7월 18 17:28:19 2024] Working: B
[목  7월 18 17:28:19 2024] Working: B
[목  7월 18 17:28:19 2024] Working: B
[목  7월 18 17:28:19 2024] Working: B
[목  7월 18 17:28:19 2024] Working: B
[목  7월 18 17:28:19 2024] Working: B
[목  7월 18 17:28:19 2024] Working: B
[목  7월 18 17:28:20 2024] Working: B
[목  7월 18 17:28:20 2024] Working: B
[목  7월 18 17:28:20 2024] Working: B
[목  7월 18 17:28:20 2024] Working: B
[목  7월 18 17:28:20 2024] Working: B
[목  7월 18 17:28:20 2024] Working: B
[목  7월 18 17:28:20 2024] Working: B
[목  7월 18 17:28:20 2024] Working: B
[목  7월 18 17:28:21 2024] Working: B
[목  7월 18 17:28:21 2024] Working: B
[목  7월 18 17:28:21 2024] Working: B
[목  7월 18 17:28:21 2024] Working: B
[목  7월 18 17:28:21 2024] Working: B
[목  7월 18 17:28:21 2024] Working: B
[목  7월 18 17:28:21 2024] Working: B
[목  7월 18 17:28:21 2024] Working: B
[목  7월 18 17:28:21 2024] Working: B
[목  7월 18 17:28:21 2024] Working: B
[목  7월 18 17:28:22 2024] Working: B
[목  7월 18 17:28:22 2024] Working: B
[목  7월 18 17:28:22 2024] Working: B
[목  7월 18 17:28:22 2024] Working: B
[목  7월 18 17:28:22 2024] Working: B
[목  7월 18 17:28:22 2024] Working: B
[목  7월 18 17:28:22 2024] Working: B
[목  7월 18 17:28:22 2024] Working: B
[목  7월 18 17:28:22 2024] Working: B
[목  7월 18 17:28:22 2024] Working: B
[목  7월 18 17:28:23 2024] Working: B
```

5.3 Analysis

In your analysis, you should compare the wait time and the response time of the processes under each of the scheduling policies. To effectively present your results, provide a graphical representation. Specifically, **draw a graph** that compares the wait times and the response time of the three processes across the different scheduling policies. Additionally, **discuss the characteristics** of each policy, such as their pros and cons, and the proper cases (situations) for the use of each policy.

6. Restrictions

1. Use the kernel log (**dmesg**) to watch the behavior of **ku_cpu**.
2. The queue in **ku_cpu** can be implemented as an array or any other structure type.
3. It is recommended to modify the source code examples in this instruction.
4. It is acceptable to have some discrepancies in the experiment results (including latency) as our project is a kind of simulation.

7. Submission guidelines

1. Submit all files through Blackboard.
2. List of files to submit
   A. Report
   B. Your source code: you should write comments for major (important) variables and functions
   C. Text file that contains execution results and log: save the log of user processes and scheduling-related kernel log
   D. **Videos (.mp4)** for executing "**run"** and observing logs: record the screen you are running the "run" process. Also, the video should include logs from the experiments. You can use a screen-capture function of your PC or any other convenient tool for the video recording.
3. How to submit files
   A. Please follow below naming rules including file extensions and directory structure.
   B. Compress all the files and directories into a "os2_[your student ID].zip".

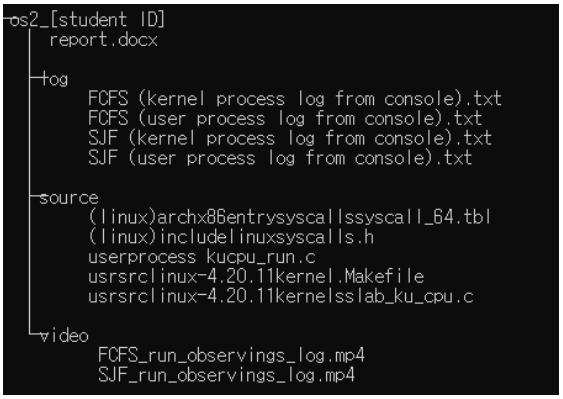

   C. **Never submit the entire Linux kernel source code.**
   D. All source codes must be able to be compiled and executed in an identical Linux environment.

4. Be sure to include the following sections in the report (The report should not exceed 11 pages, including the cover page):
   - **A. We provide a skeleton report for your convenience. Be sure to use the skeleton.**
   - **B. You can use either English or Korean to write the report.**
   - C. Please list the following on the cover: Name, student ID, department, submission date, number of free days used for Project 2
   - D. Development environment
   - E. Explanation of CPU scheduling and policies
     - i. **Caution:** Do not copy or paste existing material (including web and GitHub repositories). Although there is abundant data related to scheduling policies, ensure you understand the content and write it in your own words.
   - F. Implementation
   - G. Experiment results & analysis
   - H. Advanced question

5. Note
   - A. For late submissions that exceed the “free day” allowance (7 days in total), 1 point will be deducted from the total score (not from project 2’s score) for each day of late submission.
   - B. ***No project results will be accepted after one week (7 days)*** **has passed from the deadline of the project #2.**

8. Grading Rubric

1. Submit all files through Blackboard. Other methods will not be accepted.

2. Scoring for each submitted item

- A. Report [80pts]
  - **i. Limit your writing to 11 pages including cover page. If the format is incorrect, 1 point will be deducted from the total score. (font size: 13 pt, font type: Times new roman or 맑은 고딕)**
  - **ii. 1 point will be deducted for missing items (name, student #, Dept, date, free day, Development environment, font size, font) from the total score**
  - **iii. Explanation of CPU scheduling and policies [20 pts]**
    1. Please answer the given questions.
    2. If there is no explanation provided, a deduction of corresponding points will be applied.
    3. If the provided contents contain incorrect explanation, a deduction of corresponding points will be applied.
    4. If the explanation is insufficient, a deduction of half points will be applied.
  - **iv. Implementation [40 pts]**
    1. Zero points will be awarded if you attach screenshots(image) instead of text.
    2. If your code does not produce proper result, zero points will be awarded.
    3. For A, B, D, E parts in the report
       - A. Please attach only added or edited part of code in text. Otherwise, zero points will

be awarded.

4. For C part in the report
   A. Describe the workflow focusing on the decision made by the scheduling policy for the implemented system call.
   B. If the explanation is inadequate, a deduction of half points will be applied per each scheduling.

v. Experiment results & Analysis [10 pts]

1. Please explain each scheduling method and create a line graph.
2. If you do not include the graph, a deduction of 2 points will be applied
3. If there is no explanation provided, a deduction of 6 points will be applied.
4. If the provided contents contain incorrect explanation, a deduction of 6 points will be applied.
5. If the explanation is insufficient, a deduction of 3 points will be applied.

vi. Advanced Question [10 pts]

1. Please answer the given question.
2. If there is no answer provided, a deduction of 10 points will be applied.
3. If the provided answer contains incorrect contents, a deduction of 10 points will be applied.
4. If the answer is insufficient, a deduction of 5 points will be applied.

B. Source code [10 pts]

i. Please gather source code files (and Makefile) into a directory named "source" and there should be 5 files in the source folder

1. (linux)/arch/x86/entry/syscalls/syscall_64.tbl
2. (linux)/include/linux/syscalls.h
3. /usr/src/linux-4.20.11/kernel/sslab_ku_cpu.c
4. /usr/src/linux-4.20.11/kernel/Makefile
5. userprocess: kucpu_run.c

ii. 1 point will be deducted for incorrect file names, missing files, or submitting more than the 5 required files

iii. Write all system calls and the functions to be used in the system calls within the "sslab_ku_cpu.c" file. Points will be deducted for any other files.

C. Text file that contains execution results and log [4 pts]

i. Please gather log files into a directory named "log" and there should be 4 files in the log folder.

1. FCFS (user process log from console).txt
2. FCFS (kernel log from dmesg).txt
3. SJF (user process log from console).txt
4. SJF (kernel log from dmesg).txt

ii. 1 point will be deducted for incorrect file names, missing files, or submitting more than the 4 required files

D. Video [4 pts]

i. **Please gather video files into a directory named "video" and there should be 2 files in the video folder.**
   1. FCFS_run_observings_log.mp4

ii. **1 point will be deducted for unsuitable file name, missing file, or additional file beyond the 2 required.**

iii. **The 'user process' and 'kernel log' should be visible simultaneously. 1 point will be deducted if not visible.**

E. os2_[your student Id].zip [2 pts]

i. **Please gather files and folders into a directory named "os_[your student Id]"**

ii. **1 point will be deducted for each incorrect each directory or file name or extension.**

NOTE:

This grading rubric is to ensure fairness among students.
All grading will be conducted based on the criteria (rubric) stated above.
Most importantly, it is intended to help you complete this assignment and avoid losing scores by mistake.
If you have any questions, please feel free to contact TAs via email.