

Lab 4 : CPU Scheduling

September 6, 2016

Objective :

- Lab 4 is intended to help you learn some of the CPU scheduling algorithms discussed in class by implementing and simulating their performance.

Instructions:

- You need to program the following scheduling algorithms (Preferably in *C*): **FCFS**, **SJF (with preemption)** and **Round Robin Scheduling**
- Each program will read from a file containing a list of processes with pertinent data for the process. The program will simulate the execution of the processes. It will print out the time taken by each process to complete (*turn around time*) and the *wait time* and compute the *average turn around time* for all processes to execute as well as the standard deviation of that average.
- Inputs : A filename from the keyboard, read the file for the pertinent data, and a possible time slice size (depending on which algorithm is used). The file containing the information on the processes will have each process on a separate line. The processes will be in the file in the order in which they arrive at the OS. Each line will have a process name that will be a string. Following that will be the arrival time of the process. The arrival time will be in reference to the previous process. Following this will be the total execution time. Next will be the elapsed time between I/O interrupts (system calls), next will be the time spent waiting and processing the I/O and finally the priority of the process as an integer (smaller values will have higher priority). It will look like this

```
P1  0  20.0  1.5  5.0  2
P2  2  15.0  2.0  6.0  1
P3  6  27.0  1.8  3.5  4
p4  4  36.0  2.1  2.6  3
.    .    .    .
x   x   x    x   x
```

The xx indicates the end of the data.

- Outputs: A prompt for which file is to be read. A prompt for the amount of time for a time-slice. A list of each process and the time it took for it to complete. Then an average time for processes to complete. Finally the

standard deviation for the average time for processes to complete.

Example:

The name of the file to be read: < *filename* >

process name	turn around time	total wait time
?	?	?
?	?	?

The average time for processes to complete was ???

The standard deviation for the average process completion time was ???

- For RR Scheduling : The time slice if required for your algorithm will be 3. You may assume that a swap (context switching time) is small enough that it can be ignored. Vary the time slice/quantum of RR scheduling from 1 to 10sec(in steps of 1 sec) and plot a graph showing how the *average turnaround time* for processes vary with time slice/quantum. Also, plot a graph showing how the *average waiting time* for processes vary with time slice/quantum.