Exercise 13 for Students of Computer Science

Leon Tabak
l.tabak@ieee.org

28 September 2021

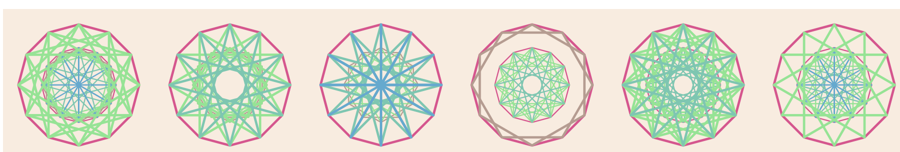

Simulation

A process is "a program in execution." A processor runs processes. An operating system schedules the execution of processes on a processor.

Write a program that simulates a machine that runs processes non-preemptively in a first-come/first served order. (Non-preemptive means that once a process begins running, it runs to completion.)

Your program will create processes and a waiting line (a queue) to hold processes.

Two non-negative floating point numbers will characterize each process. One number will be a measure of how much time it takes the processor to execute the process. The second number will be a measure of the time that elapses between the arrival of successive processes—the number on the $(n+1)^{th}$ process will signify the time that elapses between the arrival of the n^{th} process in the waiting line and the arrival of the $(n+1)^{th}$.

These two numbers will be random numbers, drawn from an exponential distribution, and having specified means. The ratio of the two means will determine the outcome of experiments carried out with your program. Random numbers drawn from an exponential disribution are non-negative, have no upper bound, and are more likely to be small than large.

Your program will compute for each process the time at which the process is first ready for execution, the time at which execution begins, and the time at which the processor completes the execution.

At the end, your program will compute the mean time that processes wait for a turn with the processor, the minimum wait for any process, and the maximum wait for any process.

You might also try tracking how the length of the waiting line grows and shrinks during an experiment.

Here are the properties of a process:

x_0 is the time that will be required to execute the process.

The value of this variable will be a random number drawn from an exponential distribution with a specified mean.

x_1 is the time that elapses between the arrival in the system of the previous process and the arrival of this process

The value of this variable will be a random number drawn from an exponential distribution with a specified mean.

x_2 is the time at which this process arrives in the system.

The program computes this value by summing that the time at which the previous process arrived and the time that elapsed between the arrival of the previous process and the arrival of this process.

If p_n and p_{n+1} are successive processes, then...

$$p_{n+1} = p_n.x_2 + p_{n+1}.x_1$$

x_3 is the time at which the processor begins executing this process.

This number is the larger of the time at which this process arrives and the time at which execution of the previous process end.

$$p_{n+1}.x_3 = \max(p_{n+1}.x_2, p_n.x_4)$$

x_4 is the time at which execution of this process ends.

This number is the sum of the time at which execution of this process begins and the time required to execute this process.

$$p_n.x_4 = p_n.x_3 + p_n.x_0$$