

**CS 115 Computer Simulation, Assignment #2 – Train Unloading Dock (again)**  
**Due at the start of the 12<sup>th</sup> lecture of the quarter.**

In this assignment, you will write a simulation of a train unloading dock. The system being modeled is *exactly* the same as that described in assignment 1, except this time you will write your simulation using the *process interaction* model, in any special-purpose simulation language or library that supports it. (Previously we used the **CSIM** language, but you can use any special-purpose system so long as it uses the process interaction model and we can run/test it on openlab; **SimPy** is likely easiest.) As with Assig. #1, it will be tested for output correctness on the ICS openlab Linux computers.

The input and output specifications remain the same; we would like to be able to run your code with various parameters and see the output. Output the same list of statistics at the end of your simulation as specified in Assignment #1. Contrary to A1, you should use **separate random number streams for each** of the random processes (train arrivals, train unloading times, remaining crew time, replacement crew travel time).

The grading guidelines (*i.e.*, “pretty” source code, correct simulation, brief but thoughtful write-up describing why you think your simulation works including various “sanity” tests) also remain the same. The late penalty is the same, and you submit both on openlab (using the Unix “submit” command as before) and on GradeScope.

In addition, you will compute two more statistics in (or at least from) your simulation(s):

- a) Compute the 99% confidence interval for the mean time-in-system (from arrival to departure) across trains, based upon 100 runs of the simulation. (Note: it’s a population average, not a time average)
- b) Compute the mean time-in-system to an accuracy of 1%, with 99% confidence. How many runs did it take to compute this? (That is, keep re-running your simulation, each time with a different seed, until your 99% confidence interval has a width which is less than 1% of the value of the mean time-in-system.)

These statistics can be computed by hand by running your simulation many times; or, if you are clever, it can all be done with a little extra coding if batch runs and confidence intervals are supported by your special purpose language/library. (In CSIM it was the functions **reset**, **permanent\_table**, **table\_mean**, **report**, **report\_table**, and the part of the CSIM User Manual discussing confidence intervals... for other systems or libraries you’re on your own). Note also that if you do it this way, you might get different numbers than if you do separate simulations and combine them by hand, because in the former case you won’t start with an empty queue at the beginning of each batch, but in the latter you will.

In addition, use your new simulation to answer the following questions:

- 1) At what average inter-arrival time does the system become overloaded? That is, leaving all distributions the same except the train arrival rate, how small can the average inter-arrival time be before the system becomes overloaded? How did you determine whether the system was overloaded or not?
- 2) Does your system overload at the same rate of train arrivals as your first assignment? If it does not, explain why not (which means that at least one of your simulations has a bug).
- 3) **(Bonus 10%)**: Provide an analytical estimate for the maximum train arrival rate (or equivalently, the minimum average inter-arrival time), if all other values and distributions remain as they are. How close does your analytical formula agree with the simulated one(s) above? If they disagree, why? Which is more accurate?