SA421 – Simulation Modeling Fall 2015
Asst. Prof. Nelson Uhan

Lesson 6. Replicating Simulations

1 Overview

- So far, we've computed and observed performance measures for 1 simulation run
  - e.g. average delay in the Fantastic Dan problem
- The observed average delay can differ between simulation runs
- Average delay is a random variable
  - Uncertain quantity before the simulation is run
  - Depends on interarrival times and service times, which are random variables
- Can we estimate the distribution of average delay?
  - Let's focus on estimating the mean and variance of this distribution

2 The experiment

- Replicate the simulation n times
- Compute performance measure (e.g. average delay) for each simulation run
  (obtaining n observations of the performance measure)
- Use the n observations to estimate the mean and variance of the performance measure

3 After the experiment: observed sample mean and sample variance

- Let $X_1, \dots, X_n$ be independent and identically distributed (i.i.d.) random variables with unknown mean μ and variance σ^2
- Let $x_1, \dots, x_n$ be the observed values of $X_1, \dots, X_n$, respectively
- For example:
  - Think of X_i as the average delay in the ith simulation run before the experiment
  - Think of x_i as the observed average delay in the ith simulation run after the experiment
  - Since the simulation runs replicate the same system, $X_1, \dots, X_n$ should be identically distributed
- We want to estimate μ

- The **observed sample mean** is

- The **observed sample variance** is

- The **observed sample standard deviation** is

- We estimate μ using the the observed sample mean
- We estimate σ^2 using the observed sample variance
- We estimate σ using the observed sample standard deviation
- These are **point estimates** for μ, σ^2 and σ, respectively
- Why should we estimate μ and σ^2 this way?

4 Before the experiment: sample mean and sample variance

- The **sample mean** is
  - The sample mean $\overline{X}$ is a random variable: before the experiment, it is an uncertain quantity
  - $\mathbb{E}[\overline{X}] = \mu$, $\text{Var}(\overline{X}) = \sigma^2/n$:

- The **sample variance** is

- The **sample standard deviation** is
  - The sample variance and sample standard deviation are also random variables: before the experiment, they are uncertain quantities

- The sample mean is an **unbiased estimator** of μ, and the sample variance is an **unbiased estimator** of σ^2: that is,

  - Intuitively, this indicates that using the observed sample mean to estimate μ and the observed sample variance to estimate σ^2 is not a bad idea

5 Confidence intervals: how good is the observed sample mean as an estimate?

- Is the observed sample mean $\overline{x}$ "close" to μ?
- Suppose $\overline{X}$ is normally distributed
  - This is true if $X_1, \dots, X_n$ are normally distributed
  - This is approximately true by the Central Limit Theorem if $n \geq 30$
- Then the $(1 - \alpha)100\%$ **confidence interval for** μ is

  - This is an **interval estimate** for μ
  - Note: more observations (larger n) $\Rightarrow$ smaller confidence intervals
  - The t-distribution with $n - 1$ degrees of freedom $\approx$ standard Normal distribution when $n \geq 30$

6 Interpreting confidence intervals

- Sample mean $\overline{X}$ and sample standard deviation S^2 are random variables
- Every experiment, we get different observed sample mean $\overline{x}$ and observed sample variance s^2

$\Rightarrow$ Every experiment, we get a different confidence interval

- After running the experiment many times, $(1 - \alpha)100\%$ of the resulting confidence intervals will contain the actual mean μ
- We say that "we are $(1 - \alpha)100\%$ confident that the mean μ lies within the confidence interval"
- Wrong interpretation: "The mean μ lies within the confidence interval with $(1 - \alpha)100\%$ probability"
- Smaller confidence interval $\Rightarrow$ more accurate estimate of μ

Example 1. Suppose an estimate of μ within 0.1 was desired at a confidence level of 95%. We perform a "warm-up" experiment of $n = 30$ simulation runs to compute an observed sample variance s^2, which is found to be 3.2. How many simulations runs are needed to obtain this estimate of μ?