Professor Abraham Matta

Basics of Performance Evaluation by Simulation

Simulation, specifically *discrete-event simulation*, models the evolution of the state of the system at discrete time instants when events occur. For example, the state of a single queue system could be defined as the number of packets in the system, and events of interest represent packet arrivals and departures. Simulation is often used to validate assumptions often made to make the analysis more tractable. However, simulation does not usually capture all the details found in actual implementations, e.g. thread scheduling overhead and bit-level transmissions, which might significantly slow down the simulation.

We review here generic simulation basics and associated statistics background.

When using simulation, we are often interested in the *steady-state average* performance of the system being simulated. To this end, given stochastic input parameters, we obtain several independent evolutions (runs) of the system starting from independent (different) random seeds (i.e. starting points for the random number generator of the input values). For each evolution (also known as *sample path*), we first ignore the so-called *transient period*, which is from time zero until the earliest point in time after which computing the time average of the performance measure (e.g. throughput) remains almost the same. Let x_i be the steady-state value of the performance measure in evolution (run) i. Then, the x_i are iid (identically independently distributed), and the sample mean for N runs is given by:

$$\overline{x} = \frac{\sum_{i=1}^{N} x_i}{N}$$

$\overline{x}$ itself is a random variable, and its distribution is that of sample means, since a different set of N runs would likely yield a different $\overline{x}$.

The question is: how close is $\overline{x}$ to the true mean of the distribution μ?

First, let us find $E(\overline{x})$, the mean of the distribution of sample means. We have:

$$N\,\overline{x} = \sum_{i=1}^{N} x_i$$
$$N\,E(\overline{x}\,) = \sum_{i=1}^{N} E(x_i) = N\,\mu$$

Thus, $E(\overline{x}) = \mu$.

Our goal is to compute a confidence interval around $\overline{x}$ within which the true mean μ lies with probability $1-\alpha$.

$$P(\overline{x} - \Delta \le \mu \le \overline{x} + \Delta) = 1 - \alpha.$$

This is can be re-written as:

$$P(\mu - \Delta \leq \overline{x} \leq \mu + \Delta) = 1 - \alpha.$$

So, another way to look at it is that our sample mean should be within a small distance from the true mean.

Δ should depend on the standard deviation of the distribution of sample means $S_{\overline{x}}$, that is, $\Delta = d\, S_{\overline{x}}$ for some factor *d*. We first find the relationship between $S_{\overline{x}}$ and the true standard deviation σ.

$$N\,\overline{x} = \sum_{i=1}^{N} x_i$$

$$N^2\, V(\overline{x}) = \sum_{i=1}^{N} V(x_i) = N\,\sigma^2$$

$$V(\overline{x}) = \sigma^2/N$$

$$S_{\overline{x}} = \sigma/\sqrt{N}$$

Notice that as *N* increases (i.e. more simulation runs), $S_{\overline{x}}$ decreases, and so Δ decreases, which means that the sample mean gets very close to the true mean.

To calculate Δ, we need to find σ. Since we don't know the true distribution, we calculate an estimate of σ as:

$$\hat{\sigma} = \sqrt{\frac{\Sigma\,(x_i - \overline{x})^2}{N-1}}$$

To compute *d*, if *N* is large (i.e. *N* > 20 typically), then the distribution of sample means approaches a Normal distribution according to the Central Limit Theorem. Thus, we take $d = z_{1-\alpha/2}$, which is equal to 1.645 for $\alpha = 0.1$.

On the other hand, if *N* is small, then the distribution of sample means is a student t-distribution. Then $d = t_{N-1,1-\alpha/2}$, which is equal to 2.132 for N=5.

It is important to show confidence intervals for any steady-state simulation results you present, and note that two graphs with overlapping confidence intervals mean that they are statistically the same.