Project #2
Simulation of an Interactive Computer System: Transient Phase Elimination and Confidence Intervals

Due: Friday, Mar. 3, 2006

This project involves the same model of an interactive, multiprogrammed computer system as used in Project 1 (fixed routing probabilities only). The purpose is to compute confidence intervals for the data points you obtained.

In order to use the technique discussed in class and in the text for computing confidence intervals, you must first eliminate the transient phase data. You must also use multiple simulation runs to generate *independent* and identically distributed random variables from which to compute the confidence intervals.

Use convergence of the cross-simulation sample means to determine when the transient phase ends. Establish a reasonable criterion for identifying "convergence." The sequence of cross-simulation sample means is likely to be somewhat "noisy;" that is, a plot of these sample means vs. the sample number (the first sample from each run, the second sample from each run, etc.) is not likely to be particularly smooth. You can improve the "smoothness" by using more values for each cross-simulation sample mean, and this means using more simulation runs.

Use enough simulation runs to make the sequence of cross-simulation sample means sufficiently "smooth" to render your convergence criterion meaningful. Experience suggests that this may be a substantial number of runs, perhaps hundreds or more, but the length of each run is probably going to be short, perhaps on the order of 500 to 1000 samples (but check this!).

After determining the end of the transient phase, collect the data you need from the steady-state phase of each of several simulation runs. *You do not need to collect steady-state data from the same simulation runs that you used to determine the end of the transient phase*. It would be more efficient to do so, but this tends to complicate the data collection enormously.

Collect enough data from each run, and use enough runs, to achieve a 90% confidence level for a confidence interval width that is no more than 10% of the value of the sample mean (remember, this is the sample mean of the sample means from each of the simulation runs). You should try to make the confidence interval width reasonably close to this 10% relative value; making it much smaller will be considered inefficient.

Turn in a project report with a cover page and three sections (plus an appendix). The cover page should state, on separate lines,

ELEC 428
Spring Semester 2006
Project 2

(date submitted)

(your name)

The body of the report should consist of:

1. An introduction that includes a summary of the project objectives.

2. A description of how you determined the end of the transient phase, how you eliminated the transient phase data from each simulation run, and how you determined when the confidence interval was small enough. Include information about how you started each of the simulation runs "randomly."

3. A presentation of the results. Specifically, show the system response times for each of the indicated number of terminals, as well as the confidence interval at each point. Present these results in two ways. First, create a table with one row for each number of terminals simulated. On each row, list

   - the number of terminals,
   - the system response time (the average total time spent by a job at the CPU and disks "per throughput"),
   - the width of the confidence interval (not the endpoints),
   - the number of simulation runs used to achieve the desired confidence interval width, and
   - the number of samples (from the steady state phase) in each simulation run.

   Also provide a plot of the system response time as a function of the number of terminals. On the same plot, show the upper and lower endpoints of the confidence interval. Put the table and plots in context with a brief and concise description of what they are.

   Plot the cross-simulation sample means as a function of the sample number, from the data you collected to determine the end of the transient phase, one plot for each number of terminals. Show on each plot where you estimate the transient phase ends.

   For this project, attach as an appendix your well-commented simulation program.

The project will be graded as follows:

Correctness:

System response times	10%
Confidence interval width	40%
Transient phase elimination	20%

Report:

Organization and completeness:	10%
Evaluation/explanation	10%
Quality*:	10%

Correctness includes response times, confidence interval widths, the end of the transient phase, and number of samples required. This is stochastic simulation, so your results will not be exactly the same as those of your fellow students. In particular, the number of samples required may vary substantially based on how each person chooses to set the initial state of each simulation run (which must be set at random).

Organization and completeness of the report: Did you follow the directions regarding the project? Are all of the sections present, clearly delineated, and complete? Did you include the table, plots, and appendix?

Evaluation/explanation: Does the text of the report provide all of the requested information? Is the explanation of the results reasonable? Did you adequately explain seemingly anomalous results? Did you fail to note suspicious or totally outrageous results? Basically, does the report demonstrate that you reflected on the results and came to justifiable conclusions about how to interpret them?

Quality relates to a number of objective and subjective aspects of the report: grammar, spelling, presentation (formatting), clarity, conciseness, and appropriateness for the intended audience. The audience you should target for this report is someone who knows what an interactive, multiprogrammed computer system is and who has a basic understanding of Yacsim, statistics, and performance evaluation in general. That is, your audience is someone like you. Think of how you would explain the model and your results to one of your classmates if she or he were not doing the same project.

All reports should be printed and spell-checked. The table and plots should be computer-generated, although you may want to annotate or elaborate the plot by hand. As in Project 1, clarity and completeness, not volume, count.

You may compare your results for this project with those obtained by other students in the class, and you may help one another with any programming or YACSIM-related problems. You may not copy another student's programs or results. With this understanding,

SIGN THE PLEDGE