

# On Credibility of Simulation Studies of Telecommunication Networks

Krzysztof Pawlikowski, Hae-Duck Joshua Jeong, and Jong-Suk Ruth Lee, University of Canterbury

## ABSTRACT

In telecommunication networks, as in many other areas of science and engineering, the proliferation of computers as research tools has resulted in the adoption of computer simulation as the most commonly used paradigm of scientific investigations. This, together with a plethora of existing simulation languages and packages, has created a popular opinion that simulation is mainly an exercise in computer programming. In new computing environments, programming can be minimized, or even fully replaced, by the manipulation of icons (representing prebuilt programming objects containing basic functional blocks of simulated systems) on a computer monitor. One can say that we have witnessed another success of modern science and technology: the emergence of wonderful and powerful tools for exploring and predicting the behavior of such complex stochastic dynamic systems as telecommunication networks. But this enthusiasm is not shared by all researchers in this area. An opinion is spreading that one cannot rely on the majority of the published results on performance evaluation studies of telecommunication networks based on stochastic simulation, since they lack credibility. Indeed, the spread of this phenomenon is so wide that one can speak about a deep crisis of credibility. In this article this claim is supported by the results of a survey of over 2200 publications on telecommunication networks in recent proceedings of IEEE INFOCOM and such journals as *IEEE Transactions on Communications*, *IEEE/ACM Transactions on Networking*, and *Performance Evaluation Journal*. Our discussion focuses on two important necessary conditions of a credible simulation study: use of appropriate pseudo-random generators of independent uniformly distributed numbers, and appropriate analysis of simulation output data. Having considered their perils and pitfalls, we formulate guidelines that, if observed, could help to ensure a basic level of credibility of simulation studies of telecommunication networks.

## INTRODUCTION

The last decade of the 20th century will be remembered as a time when computers found their place in primary schools and private homes, and became ordinary items of equipment on desks in offices and businesses. This is also a time when the computing paradigm began its drift from computer networks to network computing. There is enormous interest, in both industry and academia, in creating an AAA network: a worldwide computer network able to offer Any information service, accessible from Any place and at Any time. Before it happens, scientists and engineers will have to investigate many challenging problems of network technology and evaluate possible solutions. These research activities are certainly accelerated by achievements in the area of scientific computing, with various easy-to-use software packages specifically designed for conducting performance evaluation studies of telecommunication networks.

In the area of telecommunication networks, as in many other areas of science and engineering, the proliferation of computers as research tools has resulted in the wide adoption of computer simulation as a new paradigm of scientific investigation, in addition to the two traditional ones: theoretical studies and experimentation. Various user-friendly simulation packages offer sophisticated graphical user interfaces, animation of simulated processes, and other features. This has led to a belief that simulation is mainly an exercise in computer programming. Furthermore, this programming can be greatly simplified, since there is a plethora of simulation languages that reduce designing simulation models of telecommunication networks to placing icons (representing basic functional blocks of networks) in appropriate locations on a computer monitor, and then initiating simulation by selecting an appropriate button from a menu bar.

One can say that we have witnessed another success of modern science and technology: the emergence of wonderful and powerful tools for

*This work was partially supported by the University of Canterbury's Grant no. U6301.*

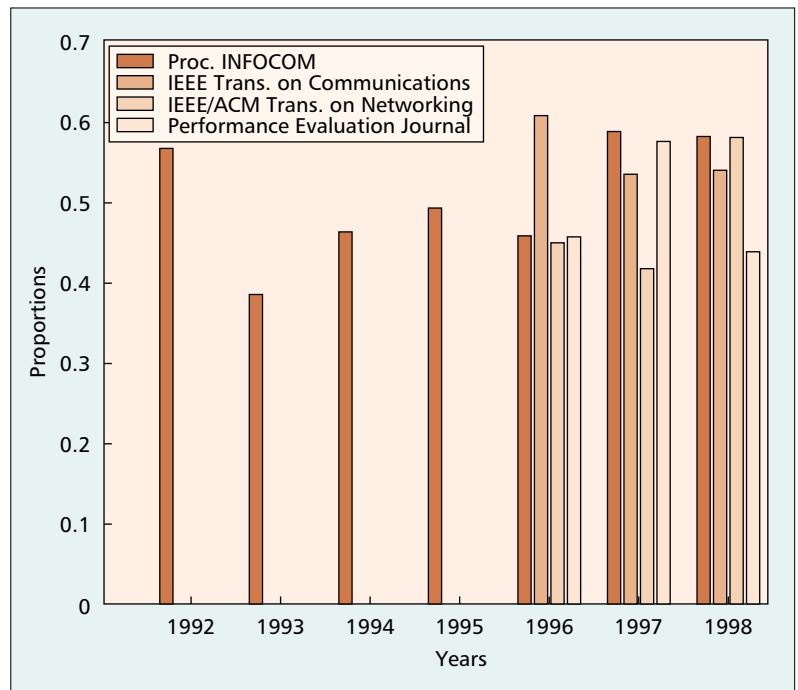
exploring and predicting the behavior of such complex stochastic dynamic systems as telecommunication networks. In fact, stochastic discrete-event simulation has already become a commonly used tool of scientists and engineers in this area, contributing to over 51 percent of all published research results (Fig. 1). The figure depicts data obtained from a survey of all papers published in proceedings of IEEE INFOCOM between 1992 and 1998 (1192 papers; ranging between 156 and 177 papers a year), *IEEE Transactions on Communications* (1996–1998; 657 papers), *IEEE/ACM Transactions on Networking* (1996–1998; 223 papers), and *Performance Evaluation Journal* (1996–1998; 174 papers). A total of 2246 papers were surveyed.

This enthusiasm is not shared by all simulation developers and users. Some claim that stochastic simulation as a performance evaluation tool of various dynamic systems, including telecommunication networks, is misused, and that the spread of this phenomenon is so wide that one can speak about a deep credibility crisis. It is even claimed that one cannot rely on the majority of the published results of performance evaluation studies of dynamic systems based on stochastic simulation. Editorials such as [1] or panel discussions of specialists on the topic, such as those organized at the Winter Simulation Conference in 1994 and 1996, or at IEEE INFOCOM in 1996, have not changed the situation.

In this article we investigate the motivation and validity of such claims. Having narrowed our interest to the application of stochastic discrete-event simulation in performance evaluation studies of telecommunication networks, we look more closely at two important necessary conditions of a credible simulation study: use of appropriate pseudo-random generators of independent uniformly distributed numbers, and appropriate analysis of simulation output data. Having considered their perils and pitfalls, we formulate guidelines that, if observed, could help ensure a basic level of credibility of simulation studies of telecommunication networks.

## THE ISSUE OF CREDIBILITY

P. J. Kiviat, in his opening address of the Summer Computer Simulation Conference (SCSC '90) [2], stated that "... *succeeding in simulation requires more than the ability to build useful models...*" Some experts assess that the modeling phase of simulation consumes only 30–40 percent of the total effort in most successful simulation projects. The first necessary step of any performance evaluation studies based on stochastic simulation is to use a *valid simulation model*. In the case of telecommunication networks, it means a valid conceptual model of the network, based on appropriate assumptions about the network's internal mechanisms, limitations, stochastic characteristics of processes that will be simulated, and so on. A good discussion of general guidelines on how to build valid simulation models can be found, for example, in [3]. However, this is only the first step in ensuring the credibility of the final results of simulation studies.



■ **Figure 1.** Proportions of papers that reported results based on stochastic simulation; obtained on the basis of the survey of 2246 papers published in Proceedings of IEEE INFOCOM, IEEE Transactions on Communications, IEEE/ACM Transactions on Networking, and Performance Evaluation Journal.

The next step is to ensure that the valid simulation model is used in a *valid simulation experiment*. At this stage, two main issues that have to be addressed for ensuring validity of any stochastic simulation-based experiment are:

- Application of appropriate elementary source(s) of randomness
- Appropriate analysis of simulation output data. Let us look closer at these two issues

### ELEMENTARY SOURCES OF RANDOMNESS

It is a generally accepted and commonly used practice today to use algorithmic generators of uniformly distributed pseudo-random numbers as elementary sources of randomness in stochastic simulation. Such an algorithm generates periodic sequences of numbers: having generated all numbers of a cycle, it repeats generation of the same sequence of numbers again. The theoretical foundations of PRNGs are well established (e.g., [4]), and over the last 50 years many different PRNGs that have passed rigorous theoretical tests have been proposed.

The most popular generators have belonged to a class of multiplicative linear congruential PRNGs (MLC-PRNGs), based on recursive algorithms in integer modulo  $M$  arithmetic. In today's world of 32-bit computers, MLC-PRNGs with the modulus  $M = 2^{31} - 1$  have been the focus of special attention, and, following exhaustive analysis, about 20 of them have been recommended as acceptable sources of independent and uniformly distributed pseudo-random numbers. These are the generators that have been used, for example, in GPSS (v. H and PC), SIMSCRIPT II.5, SIMAN, and SLAM II [3]. Thus, one would expect there to be no problem with selecting a good PRNG.

Fortunately, recent advances in the area of PRNGs have given us generators that, in a foreseeable future, should be adequate for simulations demanding even very long CPU time.

Unfortunately, this is only partially true. Conscientious users of PRNGs should be aware of potentially very serious problems with using PRNGs in real-life applications. One is that recent achievements of electronic technology have made PRNGs with cycles on the order of  $2^{32}$  obsolete in all but very short-lasting simulation studies. In 1999 a standard workstation equipped with a CPU operating with a speed of a few hundred megahertz could generate all numbers of a  $\text{mod}(2^{31} - 1)$  PRNG in 12 minutes. And this was not the final evidence of Moore's law in action: PCs with clock frequencies of over 1 GHz are already commercially available. Thus, when planning a simulation lasting longer than a few minutes of CPU time, one obviously needs PRNGs of much longer cycles than those that would have been satisfactory only a couple of years ago.

Note that simulation studies of multimedia networks (fed by streams of teletraffic modeled by strongly autocorrelated, or even selfsimilar long-range dependent, processes) can require very long simulations to obtain the final results with an acceptably small statistical error or, in other words, to collect representative samples of simulation output data (see the next section).

The use of PRNGs with adequately long cycles is also strongly advocated by recently established theoretical restrictions on the number of pseudo-random numbers from one PRNG that can be used in a single simulation. For example, if one is concerned with two-dimensional uniformity of pseudo-random numbers generated by a PRNG with cycle length  $L$ , then one should not use more than  $8^{3/2}L$  numbers from this PRNG during one simulation [5].

Fortunately, recent advances in the area of PRNGs have given us generators that, in the foreseeable future, should be adequate for simulations demanding even very long CPU time. A number of Multiple Recursive LC-PRNGs, and Combined Multiple Recursive LC-PRNGs, of cycles between  $2^{185}$  to  $2^{377}$ , can be found in [6], together with their portable implementations. Their virtual randomness has been found to be satisfactory in up to 32 dimensions.

Theoretical advances in another class of PRNGs, based on recursions in polynomial arithmetic and known as Generalized Feedback Shift Register (GFSR) PRNGs, have led to even more amazing findings. A twisted GFSR-PRNG, known as the Mersenne Twister, with a super astronomical cycle of  $2^{19937} - 1$ , and good virtual randomness in up to 623 dimensions (!), for up to 32-bit accuracy, has been proposed in [7]. Its portable implementation in C, for 32-bit machines, appears to be faster than a standard PRNG used in the ANSI C `rand()` function;<sup>1</sup> see <http://www.math.keio.ac.jp/matsumoto/emt.html> for the latest news regarding the Mersenne Twister.

Thus, at this stage, there exist PRNGs of acceptable quality, able to be used as sources of elementary randomness in stochastic simulations. However, this does not mean that all problems related with PRNGs have been solved. For example, one should be careful with using uniformly distributed pseudo-random numbers from the same generator in distributed and/or parallel simulation, because of potential correlations

existing between disjoint substreams of consecutive numbers [8]. In such types of simulation one should use PRNGs with extreme caution, since the final results can be very misleading if correlations hidden in the random numbers and in the simulated system interfere with each other.

In the case of traditional, non-distributed and non-parallel simulation on single processors, one has to be careful too. Uncontrolled distribution of various computer programs has resulted in the uncontrolled proliferation of really poor PRNGs, of clearly unsatisfactory or unknown quality. Thus, the advice given by D. E. Knuth in 1969 is even more important today, in the era of the Internet: "...replace the random generators by good ones. Try to avoid being shocked at what you find..." [4]. A longer list of useful practical guidelines on how to use, or not use, PRNGs in simulation studies can be found, for example, in [9], together with the advice that "...it is better to use an established generator that has been tested thoroughly than to invent a new one."

### SIMULATION OUTPUT DATA ANALYSIS

Any stochastic computer simulation, in which random processes are simulated, has to be regarded as a (simulated) statistical experiment and, because of that, application of statistical methods of analysis of (random) simulation output data is mandatory. Otherwise, J. Kleijnen of the University of Tilburg, the Netherlands, warns that "...computer runs yield a mass of data but this mass may turn into a mess <if the random nature of such output data is ignored, and then> ...instead of an expensive simulation model, a toss of the coin had better be used" [10]. Von Neumann, having noticed a similarity between computer simulators producing random output data and a roulette, coined the term *Monte Carlo simulation*.

Statistical error associated with the final result of any statistical experiment, or, in other words, the degree of confidence in the accuracy of a given final (point) estimate, is commonly measured by the corresponding confidence interval (CI) at a given confidence level (CL), that is, by the interval CI expected to contain an unknown value with the probability CL. In any correctly implemented simulation, the width of a CI will tend to shrink with the number of collected simulation output data, i.e. with the duration of the simulation.

Two different scenarios exist for determining the duration of a given stochastic simulation. Traditionally, the length of the simulation experiment was set as an input to the simulation programs. In such a *fixed-sample-size scenario*, where the duration of the simulation is predetermined by either the length of the total simulation time or the number of collected output data, the magnitude of the final statistical error of the results is a matter of luck. This is no longer an acceptable approach!

Modern methodology of stochastic simulation offers an attractive alternative solution, known as the *sequential scenario* of simulation or, simply, *sequential simulation*. Today, the sequential scenario is recognized as the only practical approach allowing control of the error of the final results of a stochastic simulation, since "...no procedure in which the run length is fixed before the simulation

<sup>1</sup> The ANSI C `rand()` function uses LC-PRNG in integer arithmetic with  $M = 2^{31}$ , 1103515245 as the multiplier, and 12345 as the additive constant.

begins can be relied upon to produce a confidence interval that covers <the true value> with the desired <confidence level>”[3]. Sequential simulation follows a sequence of consecutive checkpoints at which the accuracy of estimates, conveniently measured by the *relative statistical error* (defined as the ratio of the half-width of a given CI and the point estimate), is assessed. The simulation is stopped at a checkpoint at which the relative error of estimates falls below an acceptable threshold.

There is no problem with running a simulation sequentially if one is interested in the performance of a simulated network within a well-specified period of (simulated) time and the simulation output data obey the central limit theorem;<sup>2</sup> for example, when studying throughput in a network during the 24 hours of its operation. This is the so-called *terminating* or *finite time horizon simulation*. In our example, one would simply need to repeat the simulation (of the 24 hours of the network’s operations) an appropriate number of times, using different statistically independent sequences of pseudo-random numbers as sources of elementary randomness in different replications of the simulation. This ensures that the sample of collected output data (one data item per replication) can be regarded as representing independent and identically distributed random variables, and confidence intervals can be calculated using standard well-known methods of statistics.

When one is interested in studying the behavior of networks in steady state, the scenario is more complicated. First, since a steady state is theoretically reachable by a network after an infinitely long period of time, the problem lies in the execution of a *steady-state simulation* within a finite period of time. Various methods of approaching that problem, in the case of analysis of mean values, are discussed, for example, in [11]. Each of them involves some approximations. Most of them (except the so-called method of regenerative cycles) require that data collected at the beginning of a simulation, during initial warm-up periods, are not used to calculate steady-state estimates. If they are included in further analysis, they can cause a significant bias of the final results. Determination of the lengths of warm-up periods can require quite elaborate statistical techniques. When this is done, one is left with a time series of (heavily) correlated data, and the problem of assessing statistical errors of estimates calculated from such data. However, although the search for robust techniques of output data analysis for steady-state simulation continues, reasonably satisfactory implementations of basic procedures for calculating steady-state confidence intervals of, for example, mean values have already been published (e.g., [11]).

There are claims that sequential steady-state simulation and the associated problem of analysis of statistical errors can be avoided by running simulation experiments sufficiently long to ensure that any influence of the initial states of simulation becomes negligible. While such a *brute force approach* to stochastic steady-state simulation can sometimes lead to acceptable

Station	Traffic load		
	20%	60%	90%
1	0.090	0.048	0.101
2	0.059	0.047	0.090
3	0.103	0.055	0.120
4	0.131	0.057	0.101
5	0.110	0.038	0.137
6	0.131	0.049	0.128
7	0.109	0.080	0.084
8	0.069	0.056	0.119
9	0.081	0.080	0.108
10	0.135	0.076	0.157
11	0.208	0.062	0.172
12	0.186	0.091	0.212
13	0.226	0.091	0.145
14	0.314	0.143	0.106
15	*	0.102	0.202
16	*	0.145	0.203
17	*	0.211	0.243
18	*	0.248	0.430
19	*	*	*

**Table 1.** Relative errors, at 0.95 confidence level, of mean packet delays at stations of a version of a DQDB network described in [12], with 19 transmitting stations. The simulation lasted for 1,500,000 (simulated) time slots. Stars (\*) indicate cases where too few observations were collected to reduce statistical errors below 50 percent.

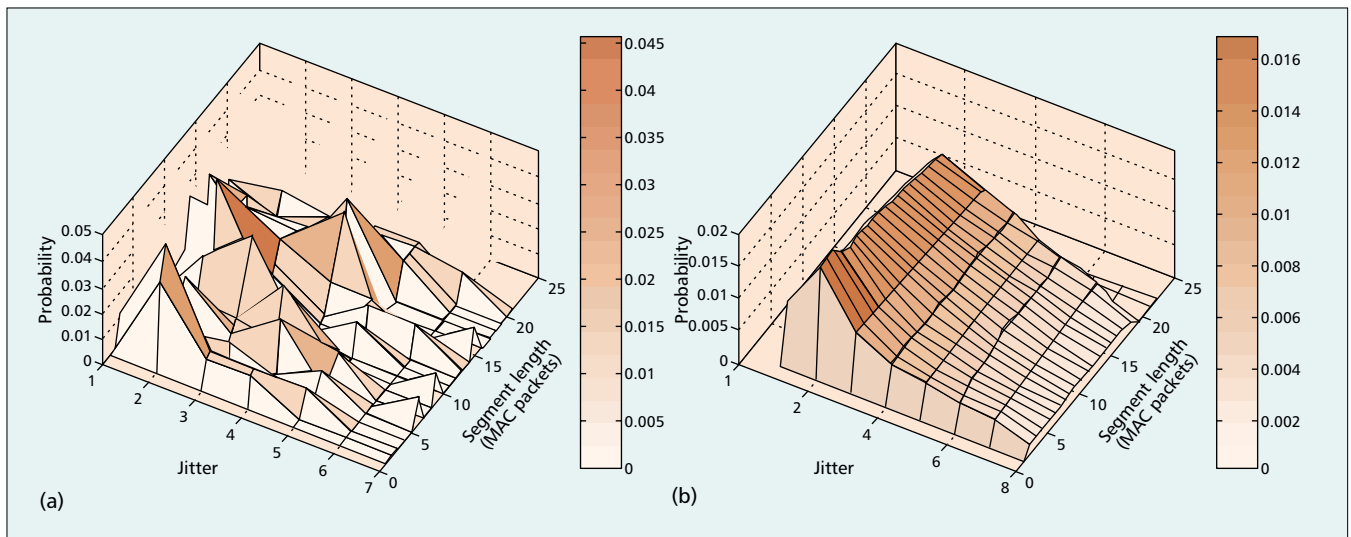
results,<sup>3</sup> one should be cautious of the statistical accuracy of the final results even in such cases.

First, such very long simulation studies could be relied on only if PRNGs of appropriate (very long) cycles were used to avoid repeating the sequence of generated pseudo-random numbers, since it can introduce superficial correlations between simulated processes. It should also be remembered that in stochastic discrete-event simulation, collecting a sufficiently large sample of data is more important than simply running the simulation over a long period. The CPU time spent on the simulation of telecommunication processes during which no event of interest is recorded does not have a direct influence on the statistical accuracy of the estimates dependent on the occurrence of these events. What matters is the number of events recorded. For example, when analyzing rare events, a minimum number of the rare events must be recorded during a given simulation to ensure that the sample of output data is representative. This phenomenon is illustrated in Table 1, which shows that estimates of the mean delays of packets obtained from a simulation (of a version of DQDB network pro-

Sequential simulation follows a sequence of consecutive checkpoints at which the accuracy of estimates, conveniently measured by the relative statistical error, is assessed. The simulation is stopped at a checkpoint at which the relative error of estimates falls below an acceptable threshold.

<sup>2</sup> The central limit theorem says that the average of a large number of output data has, in the limit (as the sample size tends to infinity), a normal probability distribution. This theorem is not applicable to, for example, output data governed by some heavy-tailed distributions.

<sup>3</sup> Some researchers execute their network simulations for a week, or longer, to get the results, they claim, representing steady-state behavior of simulated networks.



**Figure 2.** An example showing the influence of statistical errors on the quality of simulation results. The assumed confidence level = 0.9. Evaluation of a medium access control protocol of a mobile communication network proposed in [13]. a) Results with relative statistical errors of 25 percent or less; b) results with relative statistical errors of 1 percent or less.

posed in [12]) over 1,500,000 time slots can still be associated with as high a relative statistical error as 43 percent or more. The cause is clear: during this simulation many simulated time slots were idle. When there was no packet for transmission, no packet delay was measured, and no output data item was collected. Instead of stopping the simulation after a fixed time (here after 1,500,000 time slots), it should be done when the statistical errors of all estimated mean packet delays become acceptably low.

Stopping stochastic simulation too early can give misleading, or at least inconclusive, results, as illustrated in Fig. 2. The figure shows the final results obtained from sequential steady-state simulation of a wireless network operating under code-division multiple access (CDMA) (see [13] for its detailed description), but with output data analysis stopped when all estimates reached a relative statistical error not exceeding 25 percent (Fig. 2a) or 1 percent (Fig. 2b). In this context, one should certainly question the sense of drawing conclusions on the basis of results with high statistical errors, or results for which statistical errors were not measured at all.


Unfortunately, sequential stochastic simulation is not very popular among designers of commercial simulation packages, with an overwhelming majority promoting analysis of output data only after the simulation is finished. Such packages as, Arena (by Rockwell Software; see <http://www.arenasimulation.com>), Prophesy (by Abstraction Software; see <http://www.abstraction.com>), QNAP2 (by Simulog; see <http://www.simulog.fr>) or SIMSCRIPT II.5 (by CACI; see <http://www.caciasl.com>) are among the few exceptions. To this list of (commercial) packages able to execute stochastic simulation sequentially, one could also add a few packages designed at universities and offered as freeware for non-profit research organizations. One such package is Akaroa2 designed at the University of Canterbury, in Christchurch, New Zealand; see [http://www.cosc.canterbury.ac.nz/research/RG/net\\_simulation\\_group/akaroa/](http://www.cosc.canterbury.ac.nz/research/RG/net_simulation_group/akaroa/).

## CRISIS

It would probably be difficult to find a computer scientist or telecommunication engineer today who has not been trained in how to assess and minimize errors inevitably associated with statistical inference. At the same time, the results of our survey of recent publications showed that a surprisingly large proportion of papers reporting simulation-based results did not care about the random nature of output data generated by stochastic simulation. Such “don’t care” papers (represented in Fig. 3a–d by white bars) constituted 76.6, 79.05, 71.6, and 68.6 percent of papers reporting simulation-based results and published in the *Proceedings of IEEE INFOCOM*, *IEEE Transactions on Communications*, *IEEE/ACM Transactions on Networking*, and *Performance Evaluation Journal*, respectively. When the type of simulation was reported, terminating simulation (TS) and steady-state simulation (SS) appeared to be almost equally often applied in the papers published in conference proceedings, while TS dominated in journal publications. Nevertheless, the vast majority of papers did not contain any information about the time horizon over which the performance of simulated systems was analyzed (NN bars in Fig. 3a–d).

As Fig. 4 shows, on average about 76.45 percent of authors of simulation-based papers on telecommunication networks were not concerned with the random nature of the results obtained from their stochastic simulation studies, and either did not care to mention that their final results were outcomes of appropriate statistical analyses or ... reported purely random results. Let us add that Figs. 3 and 4 were obtained assuming that papers simply reporting averaged results (say, averaged over a number of replications), but without any notion of statistical errors, were included in the category of papers “with statistically analyzed results.”

While one can claim that the majority of researchers investigating performance of net-



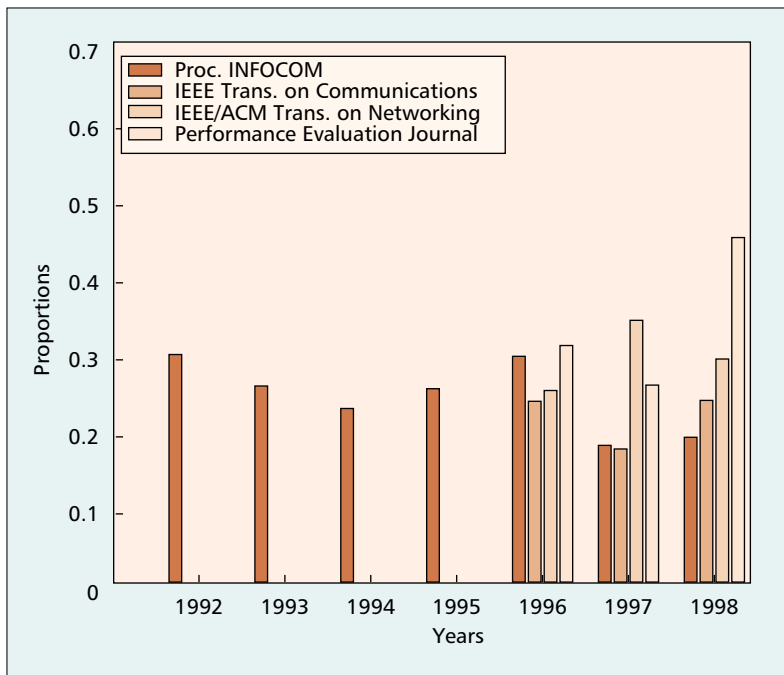
**Figure 3.** Histograms of papers reporting simulation-based results, as published in the surveyed papers in a) Proceedings of INFO-COM; b) IEEE Transactions on Communications; c) IEEE/ACM Transactions on Networking; and d) Performance Evaluation Journal. TS: papers reporting results of terminating simulation; SS: papers reporting results of steady-state simulation; NN: papers with no information about the type of simulation executed, nor output data analysis, nor statistical errors of results.

works by stochastic simulation simply may not mention that their final results have been subjected to an appropriate statistical analysis, this is not an acceptable practice. Probably everybody agrees that performance evaluation studies of telecommunication networks should be regarded as a scientific activity in which one tests hypotheses on how these complex systems would work if implemented, including even their possibly most critical conditions. However, if this is a scientific activity, one should follow the *scientific method*, a generally accepted methodological principle of modern science. This principle says that *any scientific activity should be based on controlled and independently repeatable experiments*.

Through many repetitions of a nonsequential simulation one can eventually obtain the final results with acceptably small statistical errors.

Thus, using nonsequential simulation, it is still possible to control the error of final results. However, the real problem is that the vast majority of simulation experiments reported in telecommunication network literature are not repeatable. A typical paper contains very little or no information about how a simulation was run. Our survey revealed that authors of almost 52 percent of the papers reporting simulation-based results did not even inform the reader whether their results came from terminating or steady-state simulation.

While the principles of the scientific method are generally observed by researchers conducting experiments in such natural sciences as biology, medicine, or physics, the crisis of credibility of scientific outcomes is not limited to the area of telecommunication networks but has spanned over the whole area of computer science, as well



■ **Figure 4.** Papers with statistically analyzed output data as a proportion of papers reporting simulation-based results.

as electronic and computer engineering; despite such early warnings as that in 1990 by B. Gaither, then Editor-in-Chief of *ACM Performance Evaluation Review*, who, concerned about the way in which stochastic simulation was used then, wrote that he was unaware of “...any other field of engineering or science <other than computer science and engineering> where similar liberties are taken with empirical data...” [1].

What can be done to change the attitude of writers (who, of course, are also reviewers) of papers reporting simulation studies of telecommunication networks? The consequences of drawing potentially not fully correct, or false, conclusions about network performance can be huge. Only thorough prediction of networks’ performance could ensure that telephones remain always operational and computers do not go wrong, and make such disasters as the 1990 failure of AT&T’s entire long distance network avoidable.

## A SOLUTION?

The credibility crisis of simulation studies of telecommunication networks could be resolved if some obvious guidelines for reporting the results of simulation studies were adopted.

First, the reported simulation experiments should be repeatable. This should mean that information about:

- The PRNG(s) used during the simulation
- The type of simulation

is provided. The former should be either explicitly specified in a given paper or in the documentation of the simulation package used in the study. In the latter, in the case of a terminating simulation, its time horizon should be specified.

Second, one should specify:

- The method of analysis of simulation output data

- The final statistical errors associated with the results

A satisfactory level of credibility of the final simulation results cannot be obtained without assessing their statistical errors, although sometimes, in preliminary studies, it can be acceptable to reduce the randomness of the output results of a simulation simply by repeating the simulation a number of times and averaging the results over replications. D. Knuth wrote that “...the most prudent policy for a person to follow is to run each Monte Carlo program <or stochastic simulation of a telecommunication network> at least twice, using quite different sources of pseudo-random numbers, before taking the answers of the program seriously” [4].

As mentioned, to achieve full credibility of a simulation one needs to develop valid simulation models and use them in valid simulation experiments. The former includes accurate procedural representation of the simulated system’s functionality as well as semantic and syntactical correctness of the simulation programs. The most effective way to conduct the latter is to use good, thoroughly tested PRNGs and to control statistical errors of simulation results by analyzing them sequentially; that is, to control the magnitude of the statistical errors of the results by stopping the simulation when the errors of the results reach a satisfactorily low level.

Neglecting the proper statistical analysis of simulation output data cannot be justified by the fact that some stochastic simulation studies, in particular those aimed at evaluating simulated systems in their steady state, might require sophisticated statistical techniques. On the other hand, it is true that in some cases of practical interest, appropriate statistical techniques have not yet been developed. However, if this is the case, one should not pretend that he/she is executing a precise quantitative study of the performance of a telecommunication network.

## FINAL COMMENTS

In this article we have considered two important necessary conditions of credible simulation studies: use of appropriate pseudo-random generators of independent uniformly distributed numbers, and appropriate analysis of simulation output data. Our survey of recent research publications in the area of telecommunication networks suggests that the majority of recently published results of simulation studies do not satisfy the basic criteria of credibility. The situation could easily be corrected if scientists and engineers who apply stochastic simulation as the tool for studying performance of networks accept fuller responsibility for credibility of their results. An adoption of the basic guidelines indicated in the previous section would be the first step in this direction.

Of course, simulations of telecommunication networks are often computationally intensive and can require long runs in order to obtain results at a desired level of precision. Excessive run-time hinders development and validation of simulation models. Research on speeding up the execution of simulation of telecommunication networks is one of the challenging problems

that has attracted considerable scientific interest and effort. One direction of research in this area has focused on developing methods for the concurrent execution of loosely coupled parts of large simulation models on multiprocessor computers, or multiple computers of a network. Sophisticated techniques proposed to solve this and related problems are surveyed, for example, in [14]. In addition to efficiently managing the execution of large partitioned simulation models, this approach can also offer a reasonable speedup of simulation, provided that a given simulation model is sufficiently decomposable. Unfortunately, this feature is not frequently observed in practice; thus, the efficiency of this kind of distributed simulation is strongly model-dependent.

In the context of stochastic simulation, there are still at least two additional approaches possible to speed up such simulations. One is based on variance reduction techniques (VRTs); see [3] for a survey of these. No universally applicable VRT can be found. However, such VRTs as Importance Sampling or Splitting (or Restart) can give spectacular speedup in the case of simulation of such rare events as buffer overflows or lost packets.

Stochastic simulation can also be speeded up by the concurrent generation of multiple, statistically identical streams of output data, using multiple simulation engines running statistically independent replications of simulation processes. This approach to distributed stochastic simulation is known as Multiple Replications In Parallel (MRIP) [15]. A fully automated methodology of sequential stochastic simulation in an MRIP scenario has been developed in the AKAROA project at the University of Canterbury, Christchurch, New Zealand, and implemented in Akaroa2 ([http://www.cosc.canterbury.ac.nz/research/RG/net\\_sim/simulation\\_group/akaroa](http://www.cosc.canterbury.ac.nz/research/RG/net_sim/simulation_group/akaroa)), a prototype simulation controller that, in its latest version, automatically launches multiple simulation engines and executes terminating or steady-state simulations with full control of statistical errors, in the case of such estimators as mean values, proportions, and quantiles. Further increase of the functionality of this type of simulation packages will be possible when some theoretical problems, related to the properties of new (linear combinations of) estimators, are successfully resolved.

#### ACKNOWLEDGMENTS

The authors would like to thank Don McNickle (from the University of Canterbury, New Zealand) and Frank Fitzek (from Technical University of Berlin, Germany) for their contributions in research activities that led to this article, as well as three anonymous referees for their constructive comments.

#### REFERENCES

- [1] B. Gaither, "Empty Empiricism," *ACM Perf. Eval. Rev.*, 18, 1990, no. 2, pp. 2-3.

- [2] P. J. Kiviat, "Simulation, Technology and the Decision Process," *ACM Trans. Modeling and Comp. Sim.*, 1, no. 2, 1991, pp. 89-98.
- [3] A. M. Law and W. D. Kelton, *Simulation Modeling and Analysis*, 2nd ed., New York: McGraw-Hill, 1991.
- [4] D. E. Knuth, *Art of Programming, Volume 2: Seminumerical Algorithms*, 3rd ed., Reading, MA: Addison-Wesley, 1998.
- [5] P. L'Ecuyer and R. Simard, "On the Performance of Birthday Spacing Tests with Certain Families of Random Number Generators," *Math. and Comp. in Sim.*, vol. 55, no. 1-3, 2001, pp. 131-37.
- [6] P. L'Ecuyer, "Good Parameters and Implementations for Combined Multiple Recursive Random Number Generators," *Ops. Res.*, 47, 1999, pp. 159-64.
- [7] M. Matsumoto and T. Nishimura, "Mersenne Twister: A 623-Dimensionally Equidistributed Uniform Pseudo-Random Number Generator," *ACM Trans. Modeling and Comp. Sim.*, 8, no. 1, 1998, pp. 3-30.
- [8] P. Hellekalek, "Don't Trust Parallel Monte Carlo!," *Proc. 12th Wksp. Parallel and Distrib. Sim.*, Banff, Canada, May 1998, 1998, pp. 82-89.
- [9] R. Jain, *The Art of Computer Systems Performance Analysis*, New York: Wiley, 1991.
- [10] J. P. C. Kleijnen, "The Role of Statistical Methodology in Simulation," *Methodology in Systems Modeling and Simulation*, B. P. Zeigler et al., eds. North-Holland, Amsterdam, 1979.
- [11] K. Pawlikowski, "Steady-State Simulation of Queuing Processes: A Survey of Problems and Solutions," *ACM Computing Surveys*, 2, 1990, pp. 123-70.
- [12] Q.-Y. Lee et al., "Fairness Improvements Using Multiple-Slot Reservation on QDB Networks," *Proc. SICON'91*, Singapore, Sept. 1991, pp. 79-85.
- [13] F. H. P. Fitzek et al., "Quality of Service Support for Real-Time Multimedia Applications over Wireless Links Using the Simultaneous MAC-Packet Transmission (SMPT) in a CDMA Environment," *Proc. MoMuC '98*, Berlin, Germany, Oct. 1998, pp. 367-78.
- [14] R. L. Bagrodia, "Perils and Pitfalls of Parallel Discrete Event Simulation," *Proc. WSC '96*, Coronado, CA, Dec. 1996, pp. 136-43.
- [15] K. Pawlikowski, V. Yau, and D. McNickle, "Distributed Stochastic Discrete-Event Simulation in Parallel Times Streams," *Proc. WSC '94*, Orlando, FL, Dec. 1994, pp. 723-30.

#### BIOGRAPHIES

KRZYSZTOF PAWLIKOWSKI [SM] (k.pawlikowski@cosc.canterbury.ac.nz) is an associate professor (reader) in computer science at the University of Canterbury, Christchurch, New Zealand. He received his Ph.D. in Computer Engineering from the Technical University of Gdansk, Poland. The author of over 100 research papers and four books, he has given invited lectures at about 100 universities and research institutes in Asia, Australia, Europe, and North America. His research interests include stochastic simulation, multimedia telecommunication networks, and teletraffic modeling.

HAE-DUCK JOSHUA JEONG [STM] (joshua@cosc.canterbury.ac.nz) is currently a Ph.D. student in computer science at the University of Canterbury, Christchurch, New Zealand. He received a B.Sc. in computer science and statistics and an M.Sc., computer science from Chungnam National University in Taejon, South Korea. He was a research assistant and senior programmer at the same institute from 1986 to 1993. His research interests include teletraffic modeling, stochastic simulation, and multimedia.

JONG-SUK RUTH LEE [M] (ruth@cosc.canterbury.ac.nz) received her Ph.D. in computer science from the University of Canterbury, Christchurch, New Zealand, in 2001. Earlier she received a B.Sc. in computer science and statistics and an M.Sc. in computer science from Chungnam National University, Taejon, South Korea. She was a researcher at the Electronics and Telecommunications Research Institute (ETRI), Taeduck Science Town, South Korea. Her research interests include distributed simulation and stochastic simulation of telecommunication networks.

One direction of research in this area has focused on developing methods for the concurrent execution of loosely-coupled parts of large simulation models on multi-processor computers, or multiple computers of a network.