AD-A054 945 SYRACUSE UNIV N Y F/G 9/2
LARGE SCALE INFORMATION SYSTEMS. VOLUME IV.(U)
MAR 78 F30602-74-C-0335
UNCLASSIFIED RADC-TR-78-43-VOL-4 NL

| OF |
AD
A054945

END
DATE
FILMED
7-78
DDC

RADC-TR-78-43, Volume IV (of four)
Final Technical Report
March 1978

LARGE SCALE INFORMATION SYSTEMS

Syracuse University

Approved for public release; distribution unlimited.

ROME AIR DEVELOPMENT CENTER
Air Force Systems Command
Griffiss Air Force Base, New York 13441

This report has been reviewed by the RADC Information Office (OI) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

RADC-TR-78-43, Volume IV (of four) has been reviewed and is approved for publication.

APPROVED: James L. Previtf

JAMES L. PREVITF
Project Engineer

APPROVED: Wendall C. Bauman

WENDALL C. BAUMAN, Colonel, USAF
Chief, Information Sciences Division

FOR THE COMMANDER: John P. Huss

JOHN P. HUSS
Acting Chief, Plans Office

If your address has changed or if you wish to be removed from the RADC mailing list, or if the addressee is no longer employed by your organization, please notify RADC (ISCA) Griffiss AFB NY 13441. This will assist us in maintaining a current mailing list.

Do not return this copy. Retain or destroy.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE	READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER: RADC-TR-78-43 Vol [illegible]	2. GOVT ACCESSION NO. / 3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle): LARGE SCALE INFORMATION SYSTEMS, Volume IV.	5. TYPE OF REPORT & PERIOD COVERED: Final Technical Report. 3 Jun 74 — 2 Feb 77.
	6. PERFORMING ORG. REPORT NUMBER: N/A
7. AUTHOR(s): Syracuse University	8. CONTRACT OR GRANT NUMBER(s): F30602-74-C-0335
9. PERFORMING ORGANIZATION NAME AND ADDRESS: Syracuse University, Syracuse New York 13210	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS: 62702F, 55810244, 02
11. CONTROLLING OFFICE NAME AND ADDRESS: Rome Air Development Center (ISCA), Griffiss AFB NY 13441	12. REPORT DATE: March 1978
	13. NUMBER OF PAGES: 38
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office): Same	15. SECURITY CLASS. (of this report): UNCLASSIFIED
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE: N/A

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

Same

18. SUPPLEMENTARY NOTES

RADC Project Engineer:
James L. Previte (ISCA)

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Parallel Processing
Programming Languages
Simulation
Data Base Management
Computer Architecture

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This report describes and references work conducted by Syracuse University in four broad areas: parallel processing, programming languages, modeling, and performance evaluation of generalized data management systems.

A number of applications were evaluated for processing by an associative processor architecture including air traffic control, carryless arithmetic and simulation of high-speed random logic. An extension of the typed lambda (Cont'd)

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

339600

calculus has been developed which permits the binding and application of types. User-defined types and procedural data structures are shown to be complementary tools for data abstraction. Direct and continuation semantics of the domain of flow diagrams are formulated and the properties explored. The use of transition diagrams as a tool for structured programming has been investigated. A variety of concepts and notations have been devised to facilitate reasoning about arrays.

Work relation to various simulation tasks are reported on. A tutorial on the current statistical methods of analyzing simulation output data is provided.

A number of tasks relating to file systems are discussed and a framework is advanced for describing various file organizations and operations on files.

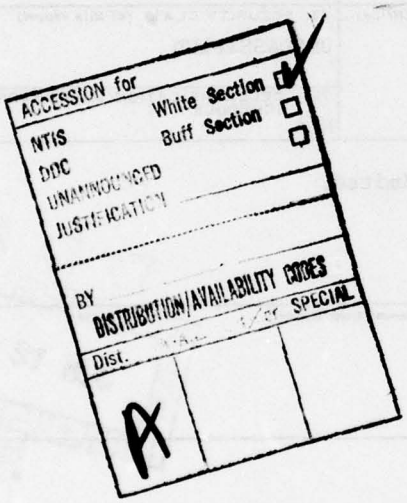

Preface

This report describes efforts completed in the Large Scale Information Systems project at Syracuse University under RADC contract F30602-74-C-0335. The work covers the period June 3, 1974 through February 2, 1977.

The report is produced in four volumes to facilitate single volume distribution.

Contents of Volume 1

Contents of Volume 2

Contents of Volume 3

Contents of Volume 4

Section 16

STATISTICAL ANALYSIS OF SIMULATION OUTPUT DATA

Robert G. Sargent

Syracuse University

Abstract

This paper is a tutorial paper on how to obtain point estimates and confidence intervals of steady state means of simulation output data. The methods of using replications, batch means, and regenerative cycles for obtaining these point and interval estimates are discussed in detail and are applied to a simple time-shared computer model to illustrate their use. A brief discussion is included on using time series methods to obtain these estimates. The advantages and disadvantages of the various methods are given, including specific recommendations as to when certain methods might be used.

TABLE OF CONTENTS

I. INTRODUCTION

The principle goal of most discrete event simulations is to investigate the steady state behavior of the model. To accomplish this goal, one of the objectives should be to determine point and interval estimates of steady state means. Unfortunately, a comprehensive statistical methodology for analyzing simulation output, including investigating the behavior of steady state means, remains to be developed. This results in many simulation studies having only point estimates of the steady state means with little or no statistical analysis of the behavior of these estimates or a specific method of analysis selected to determine the behavior of these estimates without knowing whether it is the best, what its advantages and disadvantages are, and how good the results may be. The objective of this paper is to discuss alternate methods of obtaining point estimates and confidence intervals (interval estimates) of steady state means of simulation models, including what is known about their advantages and disadvantages. We are assuming in this paper that the simulation model under analysis has a steady state mean.

After the simulation model is operational on the digital computer, the user can run the model to obtain a set of output observations. Typically the model will go through a transient state determined by the initial conditions put into the model by the user before it reaches steady state. The steady state observations are in general, not independent, but correlated, usually positively. The user analyst of a simulation model must make certain decisions to obtain point and intervals estimates of the steady state means of interest. The analyst must decide what method of data collection and data analysis to use, what initial conditions to put into the model, and what to do about the transient response, if any, under a constraint of a fixed number of observations because the amount of computer time to run the model is usually limited. There are several alternatives the analyst can choose between, but unfortunately these decisions are not as quantitive as is desired because this knowledge has

not yet been developed.

The methods of steady state analysis can be broken into two major classes. The first class, which is the simplest and the one most commonly used, is based on independent observations. One then may use the tremendous amount of statistical knowledge known regarding analyzing independent observations. This is frequently known as using classical statistics and is the type of statistics most individuals have learned. The other major class of analysis is time series analysis. This method of analysis analyzes the correlated data directly and requires knowledge of advanced statistics.

If the data generated by the simulation model is independent, then classical statistics can be used directly. However, this is usually not the case. If the analyst desires to use classical statistics, i.e. analyze independent observations, then either (1) the method of replication, (2) the method of batch means, or (3) the regenerative method, if applicable, must be selected. If time series analysis is desired, the analyst selects between spectral analysis, autocorrelation, or autoregressive-moving average methods.

The remainder of this paper is divided into six sections. The next section, Section II, discusses some of the general statistics that are used later in the paper. Section III describes three methods of obtaining independent observations and their analyses. Section IV applies the methods in Section III to a simulation model of a Time-Shared Computer System. Section V briefly discusses time series methods of analysis. Section VI discusses selection and comparisons of the various methods and the last section is the conclusion.

II. GENERAL STATISTICS

Suppose we label a set of observations from a stochastic process x_1, $x_2, \ldots, x_n$. This could be, for example, n output observations from a simulation model. A point estimator of the mean μ of this stochastic process could be obtained by using the statistic (a quantity calculated from sample observations) given by (1). If the observations $x_1, x_2, \ldots, x_n$ are indepen-

dent and identically distributed (iid), which we will assume until stated otherwise, then several additional statements can be made. First, the estimator $\bar{x}$ is an unbias estimator, i.e. $E(\bar{x}) = \mu$. Secondly, an unbias point estimator of the variance of x is given by (2). Statistics, such as $\bar{x}$ and s^2, are random variables and have distributions called sampling distributions, which describe their behavior. The variance of $\bar{x}$ is given by

$$\bar{x} = \frac{1}{n} \sum_{i=1}^{n} x_i \tag{1}$$

$$s^2 = \frac{1}{n-1} \sum_{i=1}^{n} (x_i - \bar{x})^2 \tag{2}$$

(3) and its point estimator by (4). If the variance of x is known and finite,

$$\sigma_{\bar{x}}^2 = \frac{\sigma^2}{n} \tag{3}$$

$$s_{\bar{x}}^2 = \frac{s^2}{n} \tag{4}$$

then by the central limit theorem we have (5). In general, the variance of x is not known, but it can be shown that (5) continues to hold if its esti-

$$\lim_{n\to\infty} \text{Prob}\left[\frac{\bar{x}-\mu}{\sigma_{\bar{x}}} < b\right] = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{b} e^{-y^2/2} dy \tag{5}$$

mator $s_{\bar{x}}^2$ is used to replace $\sigma_{\bar{x}}^2$. Therefore, for large sample sizes, the sampling distribution of $\bar{x}$ can be approximated by a normal distribution. One generally considers a sample size large if n is greater than thirty (30) provided that the distribution of x is reasonably well behaved.

If one desires the exact sampling distribution of $\bar{x}$, then the distribution of x must be known. For the remainder of this paper, we are going to restrict ourselves to discussing the sampling distribution of $\bar{x}$ with the variance of the random variable x unknown. If the x_i's are iid and normally distributed, the the sampling distribution of x for n observations is given by $(\bar{x}-\mu)/s_{\bar{x}}$, the t distribution with n-1 degrees of freedom. If the x_i's are not normally distributed but can be approximated by a normal, it is

common practice to use the t distribution as the sampling distribution. Unfortunately, many simulation output variables of interest have a density function that is non-negative, unimodal, and a long positive tail [2,3].

Since $\bar{x}$ is a random variable, it is usually desirable in practice to determine how accurate this point estimator is. This is usually accomplished by constructing a confidence interval for μ. The confidence interval is determined by obtaining two point estimators whose interval is called the confidence interval. This interval is chosen in such a way that one can state a probability that the interval contains the mean, a fixed but unknown quantity. In order to obtain this confidence interval, the sampling distribution of $\bar{x}$ must be known. If the sampling distribution of $\bar{x}$ is a t distribution, the $100(1-\gamma)\%$ confidence interval for μ is given by (6), where $t_{n-1,1-\gamma/2}$ is the $1-\gamma/2$ point of the t distribution with n-1 degrees of

$$\bar{x} \pm s_{\bar{x}} t_{n-1,1-\gamma/2} \tag{6}$$

freedom. Stated differently, this means that the probability that the confidence interval given by (6) contains the steady state mean is $1-\gamma$. It is also common practice to use (6) as the confidence interval of the mean if the observations are approximately normal. As n increases, the t distribution approaches the normal. For large samples, one replaces $t_{n-1,1-\gamma/2}$ in (6) by the $1-\gamma/2$ point of the standard normal, $(\mu=0,\sigma^2=1)$. A sample is generally considered large if n is greater than thirty or forty. The t distribution is a flatter distribution than the normal and its behavior is such that it gives a larger value for $t_{n-1,1-\gamma/2}$, the smaller n is. If a more accurate estimate of $\bar{x}$ is desired, then the number of observation should be increased to reduce the variance of the sampling distribution. One should also be aware that as one increases the probability that the confidence interval contains the mean, the size of the confidence interval increases.

As stated in Section I, most simulation output data goes through a transient state before reaching a steady state and the observations are

usually positively correlated. If the transient observation are included in the analysis, this will result in $\bar{x}$ being a bias estimator of μ, the steady state mean. If the data is correlated and is analyzed as it is independent, the point estimate of $\sigma^2_{\bar{x}}$ will be incorrect. Since the data is usually positively correlated, the variance of $\bar{x}$ will be underestimated resulting in a smaller confidence interval for the mean then is correct; if negatively correlated, the confidence interval is overestimated. It should be noted that most simulation languages that provide the capability of calculating variances estimators, analyze the data as if it is independent and, therefore, care should be used as to when it is appropriate to use this capability.

If the data of a simulation output is correlated and in steady state, we have what is called a covariance stationary or wide-sense stationary stochastic process. The effect of correlation does not effect estimating the mean, but does require a different method for estimating the variance of $\bar{x}$. The autocovariances of the observations are given by (7) and their point estimators by (8). The autocorrelations are given by (9) and they lie

$$R_k = E(x_t-\mu)(x_{t+k}-\mu) \tag{7}$$

$$C_k = \frac{1}{n-k} \sum_{t=1}^{n-k} (x_t-\bar{x})(x_{t+k}-\bar{x}) \tag{8}$$

between plus and minus one. Their point estimates, $\hat{\rho}(k)$, are obtained by

$$\rho(k) = \frac{R_k}{R_o} \tag{9}$$

replacing R_k by C_k. $\rho(k)$ is usually positive for simulation output data and decreases exponentially as k increases. The estimates of the variances of x and $\bar{x}$ are given by (2) and (10), respectively. One notes that if $\hat{\rho}(k)$ is zero for all k, then (10) is the same as (4). One can readily see from (10) that the variance of $\bar{x}$ is underestimated by the values associated with

$$s^2_{\bar{x}} = \frac{s^2}{n}\left[1+2\sum_{k=1}^{n-1}(1-\frac{k}{n})\beta(k)\right] \approx C_o + \sum_{k=1}^{n-1}(1-\frac{k}{n})C_k \qquad (10)$$

the autocorrelations if the simulation data is analyzed as if it is independent and is in fact positively correlated.

III. DESCRIPTION OF METHODS USING INDEPENDENT ANALYSIS

If the observations generated by the simulation model are iid, then the point and interval estimators described in Section II may be used directly for steady state means. Since simulation output observations are generally correlated, then independent observations must be obtained if classical statistics is going to be used for analysis [13, 15, 21]. There are three major approaches to obtaining independent observations and the desired estimators. They are described below.

THE METHOD OF REPLICATION

The method of replication is defined as making k independent simulation runs (replications) of length m observations each for a total of N observations. The independence of runs is accomplished by using a different stream of random numbers for each run and the same initial conditions. Let x_{ij} be the jth observation of the i^{th} run and R_i be the average of the i^{th} run. By definition the R_i's are iid. Using the R_i's in (1) and (2), the point estimators for the steady state mean μ and the variance of the R_i's are obtained, respectively. Using the s^2 calculated by (2) in (4) with n equal to k, $s^2_{\bar{x}}$ is obtained. The confidence interval of μ is obtained by using (6) with n equal to k again. Before using the R_i's described above to calculate point and interval estimates of the steady state mean, μ, one must consider two possible sources of error for these estimates.

If each run goes through a transient response before reaching the steady state response and the transient observations are included in calculating the R_i's, then $\bar{x}$ will be a bias estimator of the steady state mean μ. One must be careful of obtaining a bias in $\bar{x}$ because the confidence interval is

around $\bar{x}$, in particular, when k is large because the confidence interval decreases as k increases. It is, therefore, common practice to delete a fixed number of observations (the transient observations) from the beginning of each run before calculating the R_i's in order to eliminate the bias. However, as one eliminates observations in calculating the R_i's, the variance of $\bar{x}$ increases resulting in a larger confidence interval. This means that a trade-off must usually be made between a possible bias in $\bar{x}$ and the size of the confidence interval. The difficulty is that there is not any good procedure currently known to determine how many observations to delete.

The second source of possible error is if the distribution of the R_i's are non-normal. This would effect the sampling distribution of $\bar{x}$ and, therefore, one must be careful of using (6) in determining confidence intervals. An approach to decrease the non-normality effect is to use a large sample size (large k), bringing into effect the central limit theorem. However, this approach is usually undesirable because it increases the number of transient responses which must be eliminated to avoid a bias in $\bar{x}$. If only a total of N x_{ij} observations can be obtained, the fewer observations that must be eliminated in calculating the R_i's due to transient responses, the smaller the variance of the sampling distribution will be. This means a shorter confidence interval. One should remember that as m becomes larger, the distribution of the R_i's becomes normal.

THE METHOD OF BATCH MEANS

The method of batch means is defined as making one simulation run of length N $(x_1, x_2, \ldots, x_N)$ and dividing the set of observations from this run into k batches (segments) of m observations each. If we let B_i be the average of the m observations in batch i and choose m sufficiently large, then the B_i's should be essentially uncorrelated. If the B_i's are uncorrelated and are normally distributed, then they are also independent. Even if the B_i's are not normal, it is common practice to assume that they are

independent if they are uncorrelated. One method used to determine if the B_i's are uncorrelated is to estimate their correlations using the formulas given in Section II. Some recent empirical evidence [12, 23] has emerged that these formulas may have biases for analyzing data from simulators of queueing systems such that the estimates are lower than the actual values. If this is true then one may believe the B_i's are uncorrelated when in fact they are positively correlated. One notes that the method of batch means has only one transient response compared to the method of replication which has several transient responses.

The point estimators $\bar{x}$ and s^2 are obtained by using the B_i's in (1) and (2). The confidence interval for μ is obtained by using (4) and (6) with n equal to k. There are three possible sources of error in using the B_i's as described in calculating the point and interval estimates of the mean. If there is a transient response and the transient observations are used, then $\bar{x}$ will be a bias estimator of the steady state mean μ unless they are deleted as discussed under the method of replications. It is common practice to delete the transient response prior to dividing the run into batches for analysis.

The second source of possible error is that the distribution of the B_i's may be non-normal. This could cause the confidence interval to be the wrong size as discussed in Section II and in discussing the effects of the R_i's being non-normal. A large sample size (a large k) could also be used to bring the central limit theorem into effect for the sampling distribution of $\bar{x}$. This would not increase the number of transient responses as occurs for the method of replication.

The third possible source of error is that the B_i's may be positively correlated. If they are, we know that the variance of $\bar{x}$ will be underestimated as discussed in Section II. This would result in a smaller confidence interval than we should have. To avoid the possibility that the B_i's are uncorrelated,

the size of each batch should be as large as possible (large m). For a fixed N, a trade-off must be made between the size of k and m, which may eliminate being able to use a large sample size (large k) to have the sampling distribution of $\bar{x}$ approximately normal if the B_i's are non-normal.

THE REGENERATIVE METHOD

The regenerative method [5, 7, 8, 9, 15, 18, 19, 22] is defined as dividing a simulation run into a sequence of iid blocks by using regenerative points. If regenerative points do not exist, this method cannot be used unless approximate regenerative techniques are appropriate [10]. A model or system is said to be regenerative if there exist a sequence of points increasing in time, called regenerative points, such that the model or system begins or restarts again with the same conditions. If this method of analysis is to be applicable, the expected time between regenerative points must be finite and the model must be able to generate an infinite number of regenerative points.

The regenerative method provides a way of obtaining point and interval estimates of E(f(X)). Examples of E(f(X)) are E(X), $E(X^2)$, and P(X=0), where X can be a random variable for the steady state waiting time, steady state number is queue, steady state cost, etc. This method is more general than the two previous methods in that it can not only obtain point and interval estimators of the steady state mean, but for any E(f(X)).

Let us define a sequence of regenerative points (times) as $0 \leq \beta_1 < \beta_2 < \beta_3 < \ldots$. If the initial conditions are chosen such that they provide a regenerative point, no observations need be eliminated at the beginning of a simulation run. If the initial conditions selected are not the same as a regenerative point, then β_1 is the first regenerative point to occur in the simulation run. In some simulation models, more than one sequence of regenerative points could occur. For example, if a sequence of regenerative points occur when the model goes empty and another sequence of regenerative points occurs when

one server becomes busy, then the user could choose either sequence. Although all sets of sequences give the same length of confidence intervals asymptotically, a given simulation experiment will only have a finite number of regenerative points. The sequence chosen should be such that the number of regenerative points are reasonably large because we need a large sample size to obtain a normal distribution for the sampling distribution because it is based on the central limit theorem.

The time intervals between regenerative points are called cycles, with cycle i being the time interval between the β_i and β_{i+1} regenerative points. Let each cycle generate a Y_i and α_i, where α_i is some measure of the size of the cycle, usually either the length $\alpha_i = \beta_{i+1} - \beta_i$ or a count of the number of entities data is collected on in cycle i. Let Y_i be defined such that (11) holds, then Y_i is usually either a sum of the f(X)'s that occurred during

$$E(f(X)) = \frac{E(Y)}{E(\alpha)} \tag{11}$$

cycle i or is found by integrating f(X) over cycle i. Examples are (a) E(f(X)) = mean waiting time, then Y_i is the sum of waiting times in cycle i and α_i is the number of entities whose waiting times are in Y_i; (b) E(f(X)) = mean number in queue, then Y_i is obtained by integrating the number in queue over cycle i and α_i is the length of cycle i, and (c) E(f(X)) = probability a server is idle, the Y_i is the length of time the server is idle during cycle i and α_i is the length of cycle i.

For n cycles, a set of observations, $Y_1, Y_2, \ldots, Y_n$ and $\alpha_1, \alpha_2, \ldots, \alpha_n$ are obtained. Since the cycles are iid, the Y_i's are iid, the α_i's are iid, and the Y_i's usually highly correlated with the α_i's. There are several point and interval estimators that can be used for E(f(X)). They are a function of statistics (12) through (16). The classical point estimator of E(f(X)) is given by (17) and is a bias estimator. If we define the $100(1-\gamma)\%$ confidence interval of E(f(X)) by $r_c \pm d_c$, then the classical point estimator, d_c, is given by (18), assuming a large sample size such that sampling distribution

$$\bar{Y} = \frac{1}{n} \sum_{i=1}^{n} Y_i \tag{12}$$

$$\bar{\alpha} = \frac{1}{n} \sum_{i=1}^{n} \alpha_i \tag{13}$$

$$s_{11}^2 = \frac{1}{n-1} \sum_{i=1}^{n} (Y_i - \bar{Y})^2 \tag{14}$$

$$s_{22}^2 = \frac{1}{n-1} \sum_{i=1}^{n} (\alpha_i - \bar{\alpha})^2 \tag{15}$$

$$s_{12}^2 = \frac{1}{n-1} \sum_{i=1}^{n} (Y_i - \bar{Y})(\alpha_i - \bar{\alpha}) \tag{16}$$

can be approximated by a normal by the use of the central limit theorem. In (18), $Z_{1-\gamma/2}$ is the $1-\gamma/2$ point of the standard normal.

$$r_c = \frac{\bar{Y}}{\bar{\alpha}} \tag{17}$$

$$d_c = Z_{1-\gamma/2}\, s_c/(\bar{\alpha}\sqrt{n}) \tag{18}$$

where

$$s_c^2 = s_{11}^2 - 2r_c\, s_{12}^2 + r_c^2\, s_{22}^2$$

The Jackknife method gives better point and interval estimators of E(f(X)) than the classical method, in particular, for small sample sizes. The jackknife point estimator is given by (19) and using d_j from (20), the confidence interval for μ is $r_j \pm d_j$. Iglehart discusses the classical and Jackknife estima-

$$r_j = \frac{1}{n} \sum_{i=1}^{n} \theta_i \text{ where } \theta_i = n\frac{\bar{Y}}{\bar{\alpha}} - (n-1)\left(\sum_{j\neq i} Y_j / \sum_{j\neq i} \alpha_j\right) \tag{19}$$

$$d_j = Z_{1-\gamma/2}\, s_j/\sqrt{n} \text{ where } s_j^2 = \frac{1}{n-1} \sum_{i=1}^{n} (\theta_i - r_j)^2 \tag{20}$$

tors presented here as well as others in [18]. One should note that both the classical and Jackknife methods give better results [18] than the estimators given in the original Crane and Iglehart papers [7, 8, 9, 10] and which have been widely disseminated.

There is one possible source of error in using the regenerative methods described. The sampling distributions used in estimating E(f(X)) was obtained by using the central limit theorem. Therefore, if a large sample is not used, the sampling distribution may not be a normal. Another possible source of error could arise if one chooses to use some approximate regenerative method, [10]. One should note that the transient response is not a concern with this method. If one chooses to run a simulation for a fixed period of time, it may not end at a cycle. In this case, one uses the number of full cycles generated for data analysis and neglects the portion of a cycle left at the end of the run.

IV. A CASE STUDY

Let us apply the data analysis techniques presented in Section III to a simple time-shared computer model [1]. The model will consist of NT terminals, one CPU with round robin scheduling (a single queue with first come first serve discipline), iid think times that have an exponential distribution with a mean equal to $1/\lambda_1$, service time requests that are iid with an exponential distribution and a mean equal to $1/\lambda_2$, a maximum service quantum of length q (not including overhead), and a fixed overhead of length τ for every time slice independent of the length of the quantum. This model is given in Figure 1. Let the objective of the simulation study be to determine the point and interval estimates of the mean steady state response time, where response time, RT, is defined as the time from when a job leaves the terminal until it returns (the time between think times). The expected response time for this model can be calculated as in [1]. The model was programmed in Simscript II.5 and run on an IBM computer.

INITIAL CONDITIONS AND TRANSIENT RESPONSES

In order to simulate this time-shared computer model, a set of initial conditions must be selected. When one is interested in steady state behavior, it is desirable to select the initial state as a typical state in

steady state to reduce the transient response [6]. We are going to select a state that may not always be a typical state in steady state (it may not be for a heavy load on the model). The initial state or conditions selected is when a job leaves the computer and it goes idle. This means all jobs are in the think state. We can also use this state to generate regenerative points because think times are exponential and the exponential distribution has the forgetfulness (Markov) property.

As stated in the previous section, there does not exist any good quantitative procedure to determine when a model's transient response ends and the steady state response begins. Most simulation users simply make an "educated guess" considering such factors as (1) analysis of some realizations (replications) of interest, (2) cost of obtaining observations, (3) concern for transient effect, (4) variability of the model's behavior, and (5) congestion in the model.

This author likes to make three replications and plot the accumulative moving average of the output of interest and use these in conjunction with the factors given above in making his "educated guess" when the steady state response begins. Typically the output will either overshoot the mean value and have a damped oscillation or simply converge on the steady state value. Both of these types of response will have random fluctuations in them. The objective of using three runs is to be able to evaluate the randomness of the output between runs and to determine what fluctuations mean in a given run. One must remember in anlyzing the outputs that you are observing an accumulative moving average. Two difficulties in making these runs are what length should they be and at what interval should the accumulative moving average be printed out. The run length can always be continued if appropriate information is kept on the state of the model to be able to continue the simulation run. The interval chosen has to be large enough such that some of the randomness in the output is smooth, yet frequently

enough to observe the mean behavior.

Figure 2 contains two sets of three runs of the time sharing model, one under medium load (NT=25) and the other under a heavy load (NT=35). These runs are longer than necessary to illustrate the behavior of realizations or a simulation model. The accumulative moving average was printed out every 20 observations. Two different streams of random numbers are used in each run, one for the think times and the other for service times. The same streams are used in the model for the transient response labeled with the same number. For NT equal to 35, one can readily see that heavy congestion (heavy load) in a model causes more variability and a longer transient response than lower congestion. Some of these realizations illustrate how a transient response overshoots the mean and damps out and others illustrate converging to the mean. This author would consider the transient response ended between 200 and 300 observations for NT equal to 25 and between 600 and 700 observations for NT equal to 35.

STEADY STATE

Let us determine point and interval estimators of the steady state mean of the time-shared computer model with $1/\lambda_1 = 25$, $1/\lambda_2 = 0.8$, $\tau = 0.015$ $q = 0.1$ and $NT = 25$, with the units being seconds, using the three methods presented in Section III. This set of parameter values give a steady state mean of 3.415 seconds and is the medium load case discussed under initial conditions and transient responses above. We will use the same initial conditions given above for all our steady state investigations.

For each of the three methods, results will be presented using two replications of the experiment to illustrate the variability that can occur between experiments. For the Batch and the Regenerative Methods, the streams of random numbers used to generate realizations (1) and (2) in Figure 2 are used for the experiments (runs) one and two, respectively. The output data from each experiment is analyzed in various ways by varying

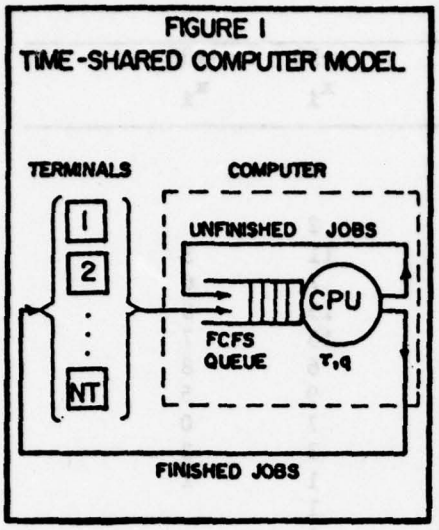

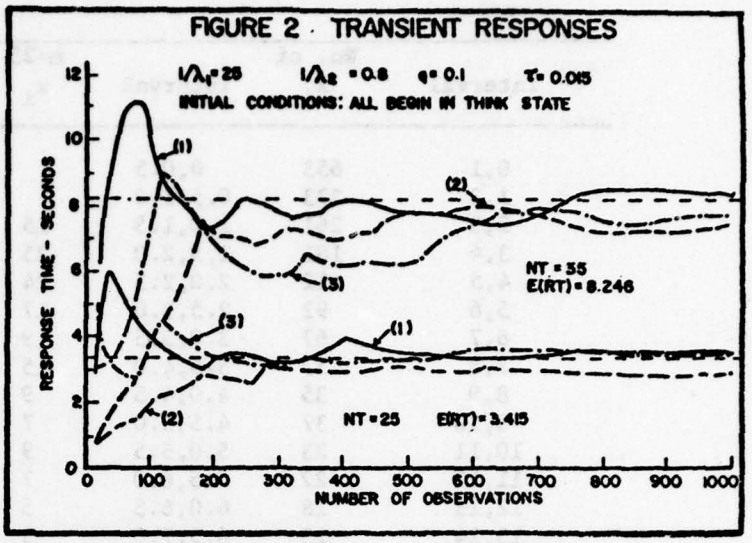

k,m,N, and the number of observations deleted for the transient response to illustrate how the results can be different depending how the analysis is done. One must not draw general conclusions from the data presented because of the variability that occurs between experiments, the differences that occur in a model's behavior for different loads, and the differences that occur between different models. The two experiments performed are actually what did happen in performing two experiments (they were not selected from several experiments). One can readily see from Figure 2 that the accumulative moving average of realization (2) never once became equal to or greater than the steady state mean in 1,000 observations. Realization (1) converges very close to the expected value during 1,000 observations.

The distribution of response time in steady state will be non-negative, unimodal, and have a long positive tail, i.e. positive skewness (skewness is a measure of asymmetry of the distribution about the mean and it is positive because the tail is in the positive direction). To illustrate the behavior of the response time distribution a frequency distribution of the data in realization (1) discussed above, is given in Table 1. One can readily see that the distribution has a long positive tail and, therefore, does not have a normal distribution.

TABLE 1. FREQUENCY DISTRIBUTIONS

Interval	No. of x_i	Interval	m=25 x_i	m=50 x_i	m=100 x_i
0,1	653	0,0.5			
1,2	373	0.5,1.0			
2,3	261	1.0,1.5	15	2	
3,4	181	1,5,2.0	25	11	3
4,5	112	2.0,2.5	24	7	6
5,6	92	2.5,3.0	17	15	8
6,7	67	3.0,3.5	19	15	7
7,8	59	3.5,4.0	15	6	8
8,9	35	4.0,4.5	9	9	5
9,10	37	4.5,5.0	7	7	0
10,11	23	5.0,5.5	9	3	3
11,12	22	5.5,6.0	7	1	1
12,13	18	6.0,6.5	5	1	
13,14	13	6.5,7.0	3	1	
14,15	12	7.0,7.5	0	1	
15,∞	42	7.5,∞	5	1	
Total No. of Obs.	2,000		160	80	40
*α_3, Coef. of Skewness			1.197	1.005	0.582

*$\alpha_3 = E(x-\mu)^3/\sigma^3$ (normal, $\alpha_3 = 0$; exponential, $\alpha_3 = 2$)

Our methods of analysis uses averages of observations and it is the distribution of these averages that must be normal to obtain an exact or approximate sampling distribution. If the observations used to calculate the averages are normal, then the averages will be normal. If the observations are not normal, the distribution of the averages will approach a normal as the number of observations in the averages increase. (Note that the central limit theorem cannot be applied directly because our observations are correlated.) Table 1 contains frequency distributions for 25, 50 and 100 observations in the averages for realization (1). One observes that as the number of observations in the averages (block size) increases, the distribution behaves more like a normal. An estimate of the coefficient of skewness for each of the three distributions are included in Table 1.

This quantifies the effect of increasing m.

If observations of response time, x_i, are collected from the model one after another, the observations will be positively correlated. Table 2 gives estimates of these correlations using the formulas given in Section II on realization (1) after removing the first 200 observations. The correlation decreases as the lag increases and one could estimate that the correlation is zero after a lag of fifteen for this set of data. In using the method of batch means, we must select the number of observations in each batch large enough such that the batch means are independent. This is determined by looking at the correlation of the batch means. If we use a batch size of 20 observations, correlation estimates using realization (1) are given in Table 2. There is some positive correlation for lag one but is zero for larger lags. This indicates we would underestimate the variance of $\bar{x}$ if we used 20 observations in each block. Using 40 observations in a block, we obtain an estimate of the first lag to be 0.011 for realization (1) and, therefore, the blocks are uncorrelated for blocks of size 40. This analysis indicates that the block size should be greater than 20 but does not need to be greater than 40. Further analysis can be done, if desired.

TABLE 2. CORRELATIONS

Response Times

t	$\beta(t)$	t	$\beta(t)$	t	$\beta(t)$
1	0.395	6	0.184	11	0.095
2	0.333	7	0.246	12	0.085
3	0.233	8	0.196	13	0.063
4	0.243	9	0.115	14	0.039
5	0.200	10	0.128	15	0.054

Average of 20 Response Times

t	1	2	3
$\beta(t)$	0.164	-0.059	0.034

Table 3 contains the results of using realizations (1) and (2) for experiements (runs) one and two, respectively, for various values of N, the total number of observations, for k, the number of batches, for m, the size of each batch, and for different number of observation deleted at the beginning of a run (transient observations). The data for N equal to 1,000 are the first 1,000 observations of N equal to 2,000, etc., for each of the runs. The estimates are obtained using the formulas in Section II as explained in Section III. $\bar{x}$ is the point estimate obtained for the confidence interval $\bar{x} \pm d$ by using formula (6). s_B^2 is the point estimate of the variance of the batches. The size of each block, m, is determined by subtracting the number of observations deleted (No. Del.) from the beginning of each run from the total number of observations, N, and dividing the difference by the number of batches, k. For example, with N equal to 1,000, k equal to 5, and the number deleted equal to zero, m equals (1,000-0)/5 = 200 observations. The estimate of the mean is the same for each run as the value obtained is not affected by the trade-off between m and k. Those confidence intervals that do not contain the mean is indicated by an * on d. d is calculated for a 90% confidence interval.

The results in Table 3 show what one would, in general, expect. $\bar{x}$, the estimate of the steady state mean, approaches E(X) = 3.415 as N gets larger. Deleting observations at the beginning of the run increases $\bar{x}$ slightly, but does not have much effect for the load and the model being analyzed. The difference between $\bar{x}$ in the two runs is noticeable, illustrating the variability between experiments. The variance of the batch averages should decrease as the size of m increases and the values of s_B^2 do. The values of d for different values of k, m, etc. is determined by $t_{1-\gamma/2,k-1}$ times $s_{\bar{x}}$, a function of k and s_B^2. The ratio of $t_{.95,4}$ to $t_{.95,39}$ is 1.265 and, therefore, the value of t does not have a significant effect. The major

effect is what happens to $\sigma^2_{\bar{x}}$ for a fixed N in the trade-off between k and m. This effect is a problem dependent. The results in Table 3 show that for N equal to 1,000 and 2,000, d increases as k decreases and for N equal to 4,000, d remains about the same. As N increases, d decreases for a given k, as expected.

These results do not determine how accurate the confidence limits are, where accuracy is defined as the 90% confidence interval actually contains 90% of the sampling distribution. One notes that three of the confidence intervals do not contain the mean and this occurs for experiment two when N equals 1,000. This occurs because $\bar{x}$ is a low estimate of μ and the confidence interval is not large enough to contain the mean. For N equal to 4,000, d is less than 10% of μ, the steady state mean.

THE REGENERATIVE METHOD

Table 4 contains the results of analyzing realizations (1) and (2) by the regenerative method using the classical estimators for various numbers of cycles. The number of observations that occurred in each cycle for each experiment is also given in Table 4 to enable a comparison between the regenerative method and the batch and replication methods of analysis whose results are given by the number of observations.

In analyzing the data in Table 4, one observes that r_c converges to the steady state mean and d_c becomes smaller as the number of cycles and observations increase, as is to be expected. The accuracy of the confidence intervals cannot be determined from the data presented. A 90% confidence interval was used. In experiment two, six of the confidence intervals do not contain the mean, including one having 226 cycles with 975 observations. In analyzing these two realizations (data not presented), most cycles have a small number of observations and a few cycles a larger number of observations. This means that there will be variability and these two experiments illustrate that. One notes that the size of the d's are approximately 10% of μ for

TABLE 3. BATCH METHOD DATA

N	No. Del.	k / Exp.	5		10		20		40		$\bar{x}$	
			1	2	1	2	1	2	1	2	1	2
1000	0	s_B^2	1.026	0.319	1.575	0.517	2.421	1.150	3.785	1.791	3.597	2.943
		d	0.966	0.538	0.727	0.417*	0.602	0.414*	0.518	0.357*		
	200	s_B^2	0.167	0.214	1.086	0.631	3.385	1.489	4.656	2.881	3.679	2.977
		d	0.390	0.441	0.604	0.460	0.711	0.472	0.575	0.452		
2000	0	s_B^2	0.333	0.343	0.752	0.441	1.409	0.943	2.080	1.677	3.298	3.344
		d	0.550	0.559	0.503	0.385	0.459	0.375	0.384	0.345		
	200	s_B^2	0.307	0.376	0.486	0.391	0.929	1.028	2.093	2.430	3.301	3.404
		d	0.529	0.585	0.404	0.362	0.373	0.392	0.385	0.415		
4000	0	s_B^2	0.083		0.209		0.471		1.017		3.377	
		d	0.274		0.265		0.265		0.269			

TABLE 4. REGENERATIVE METHOD DATA

Experiment 1				Experiment 2			
No. of Cycles	No. of Obs.	r_c	d_c	No. of Cycles	No. of Obs.	r_c	d_c
5	72	4.883	1.588	5	10	0.985	0.448*
10	78	4.546	1.588	10	19	0.851	0.302*
20	128	3.582	1.469	20	39	1.120	0.521*
40	173	3.148	1.228	40	181	2.896	0.851
80	336	3.404	0.876	80	363	2.905	0.747
160	780	3.681	0.648	160	728	2.938	0.469*
164	784	3.669	0.647	191	793	2.817	0.450*
219	988	3.604	0.561	226	975	2.887	0.386*
320	1451	3.496	0.453	320	1431	3.126	0.351
361	1584	3.378	0.433	351	1598	3.208	0.365
402	1796	3.355	0.395	396	1800	3.264	0.353
450	1988	3.301	0.367	418	1987	3.333	0.341

*Confidence Interval does not contain the mean, 3.415.

approximately 2,000 observations.

THE METHOD OF REPLICATION

Table 5 contains the results of two experiments using the method of replications for analysis. For each of the two experiments performed, forty runs were made. In the analysis of k equal to forty (forty runs), twenty of these runs were used in the analysis of k equal to twenty, etc. The length of the various runs were such to satisfy the total experiment. When N is fixed and for a given k, the size of m, the number of observations used to determine each R_i, is found by dividing N by k and subtracting the number of transient observations deleted (No. Del.). For example, if N is equal to 1,000, k equal to 5, and No. Del. equal to 50, then the m used in calculating the R_i's is equal to (1,000/5) - 50 = 150. When N is not given in the table, m is and, therefore, N can be calculated using the definitions given. Note that m is different for the two experiments when N is not specified. A 90% confidence interval has been used.

Analyzing the results in Table 5, one readily draws the conclusion that $\bar{x}$ is extremely variable. Note that a different $\bar{x}$ must be calculated for every case here. When deletion occurs for the transient responses, far less observations are available for analysis when N is fixed and this increases the variability of $\bar{x}$. When N and the number of observations deleted are fixed, s_R^2 decreases as k decreases (m increasing) and this is as expected. It is difficult to draw any conclusions about the d's except that the more observations included in the analysis, the smaller the d's become. In experiment one, a considerable portion of the confidence intervals did not contain the steady state mean. This occurs because $\bar{x}$ is very low and d is not large enough to include the value of the mean. Again, we cannot comment upon the accuracy here unless numerous experiments are performed. The sizes of the d's are much larger in Table 5 for a fixed N than they are for either the regenerative method or the method of batch means.

TABLE 5. REPLICATION METHOD DATA

N	No. Del.	k	5		10		20		40	
		Exp.	1	2	1	2	1	2	1	2
1000	0	$\bar{x}$	2.634	3.423	2.857	3.597	2.944	2.975	2.646	2.606
		s_R^2	0.237	0.595	0.437	2.785	1.161	3.276	1.994	2.066
		d	0.464*	0.736	0.383*	0.968	0.490*	0.700	0.375*	0.383*
	50	$\bar{x}$	2.672	3.659	2.962	3.939				
		s_R^2	0.127	0.641	1.371	1.825				
		d	0.340*	0.764	0.679	0.783				
2000	0	$\bar{x}$	3.228	3.493	2.798	3.538	2.934	3.352	2.881	3.068
		s_R^2	0.355	0.237	0.334	0.825	0.474	2.346	1.637	2.150
		d	0.568	0.465	0.334*	0.526	0.266*	0.592	0.341*	0.391
	50	$\bar{x}$	3.330	3.605	2.813	3.633	3.125	3.728		
		s_R^2	0.317	0.284	0.194	0.445	1.482	2.160		
		d	0.537	0.509	0.255*	0.387	0.471	0.568		
	100	$\bar{x}$	3.408	3.555	2.738	3.480				
		s_R^2	0.524	0.323	0.956	0.596				
		d	0.690	0.542	0.567*	0.448				
	200	$\bar{x}$	3.823	3.564						
		s_R^2	1.141	0.406						
		d	1.019	0.608						
	200	m	160	360	80	180	40	90		
		$\bar{x}$	3.674	3.442	3.690	3.284	2.782	3.686		
		s_R^2	1.341	0.060	1.702	0.432	1.115	1.274		
		d	1.104	0.233	0.756	0.381	0.408*	0.436		

*Confidence Interval does not contain the mean, 3.415.

V. TIME SERIES ANALYSIS

Simulation output data that is correlated and in steady state is a realization of a wide sense stationary stochastic process and can be analyzed directly using time series analysis. There are two different sets of techniques in time series analysis. One set analyzes the data in the time domain and the other in the frequency domain. Techniques that use autocorrelation, autoregressive, and autoregressive-integrated moving averages are examples of time domain techniques. Spectral analysis

[14, 17, 20] is the major technique used in the frequency domain. A large number of observations (one long run), knowledge of advance statistics, and in depth knowledge in the time series analysis technique to be used is required to apply time series analysis techniques because they are sensitive to proper application. Two advantages of using time series analysis instead of classical ststistical analysis are the correlation structure of the data is utilized in time series analysis and the simulation user is not forced to obtain independent observations.

The point estimator of the steady state mean using time series analysis is still given by (1), but estimates of the variances of the point estimator can be obtained in a variety of ways for each of the various techniques. One method of estimation for the autocorrelation technique is given in Section II. This method of estimation is not considered a good method by several investigators because the $\rho(k)$'s are correlated, the correlation estimator given has a large variance, and some empirical evidence indicate that they are bias on the low side [12,23]. Discussion of various time domain techniques are found in [5, 11, 13, 15, 21].

Spectral analysis converts data from the time domain into the frequency domain using components (sine and cosine waves) of different frequencies and amplitudes. The same information is contained in the data but different insights into the behavior of the data can be obtained by using spectral analysis as well as point and interval estimates. Data for spectral analysis is taken at fixed intervals of time. There are numerous ways to obtain estimates of the parameters in spectral analysis and they must be used with care, There has been mixed reaction in using spectral analysis for analyzing simulation output data [12, 14, 15, 16, 20].

VI. SELECTION AND COMPARISON OF METHODS

When a simulation user desires to obtain point and interval estimates of the steady state mean, one specific method of analysis must be selected.

The method selected, no doubt, will depend upon the specific problem and on the user's knowledge of analysis techniques. This author recommends that simulation users do not use time series analysis techniques to analyze simulation output unless one is very competent in using time series analysis techniques for analyzing correlated data. Most simulation users, therefore, must select between techniques using independent analysis. Unfortunately, little quantitative knowledge is available to aid in this decision.

This author recommends that a simulation user's first choice should be the regenerative method, <u>if</u> it can be used. This method has the advantage of not having to be concerned about a transient response. A possible disadvantage is that a large number of cycles must be obtained to have a normal distribution by the central limit theorem. If the expected time between cycles is large, then it may be expensive or even prohibitive to obtain a large number of them. Unfortunately, a large portion of simulations performed do not have regenerative points. Although approximate regenerative methods have been put forth, this author believes that they still are experimental, and, therefore, should not be used until more experience has been gained in their use. This results in the regenerative method not being applicable for use in most simulations. Furthermore, simulation analysts must have a good understanding of what regenerative points are.

In both the methods of batch means and replications, the transient problem must be faced. What initial conditions should be selected and how many observations should be deleted from each run? These two decisions are more important for the method of replication because they affect each replication. The more observations that are deleted, the less observations that are left for analysis. Restricting our discussion to the method of replication for the moment, this says k, the number of replications, should be small to avoid deleting as many observations as possible. For a fixed

N, this means that m will be larger. The larger m is, the smaller the variance of the R_i's will be and the closer the distribution of the R_i's will approximate the normal if the observations are non-normal. However, a large k would allow the central limit theorem to be applied to the distribution of the R_i's. The effect of the trade-off between m and k on the size of the confidence interval is problem dependent. Investigations by Law [23] shows that for simple queueing models, a small k is desirable when accuracy is the criteria (say 5 or 6).

This author recommends using the method of batch means over the method of replication because only one run is used and, therefore, there is only one transient response. The user can afford to eliminate a few extra observations to be certain the transient response has been eliminated because this is done only once. An analyst still must make a decision regarding k and m. This author believes a small k should be used to enable m to be as large as possible for a fixed N because (1) the larger m is, the less likely that the B_i's will be correlated; (2) the larger m is, the closer the distribution of the B_i's will be to the normal if the observations used for the B_i's are non-normal; and (3) the variance of the B_i's will be smaller. This means that the sampling distribution can be estimated more accurately resulting in a better accuracy for the confidence interval. The disadvantage is that a larger k may give a smaller confidence interval. Law's [23] also recommends a small k (5 or 6) based on studying simple queueing models and using accuracy as a criteria. The method of batch means, if properly used, should give results similar to the regenerative method because they both are based upon using one long run to generate iid observations.

This author suggests as an alternative to the methods of replication and block means, to combine both methods to form another method of analysis. If the number of replications is kept small, the number of transient observations deleted will not be large. Each run can be divided into a small

number of blocks of equal size. Let B_{ij} be the j^{th} block average from the ith run. The B_{ij}'s are iid if all the blocks are the same size and, therefore, the analysis is the same as for the method of block means, B_i's. This method will combine the advantages of both methods and reduce the effect of their individual disadvantages.

VII. CONCLUSIONS

This paper presented in detail three methods of obtaining point and interval estimates of the steady state mean of simulation output data and briefly discussed others. It was recommended that the regenerative method be the first choice of analysis, if applicable. The second choice recommended was the method of batch means. A case study was included to illustrate the use of the three methods presented. The results of the study showed that there can be a large variability in obtaining estimates and, therefore, care must be used in analyzing simulation output data. Caution should be used in drawing conclusions of the case study to use in other simulation studies.

Variance reduction techniques were not discussed, details of time series analysis techniques were not presented, nor were methods for comparison of alternatives. Methods to determine sample size for desired confidence intervals were not discussed. To obtain sample sizes, one must perform a simulation experiment to obtain point estimates of μ and $\sigma_{\bar{x}}^2$, then use these estimates to obtain an estimate of the sample size. See [15, 16] for details.

One must always remember that statistical analysis of the output data is only one step in the simulation methodology. The results of the output analysis is no better than how good the input data is or how valid the model is. Research is continuing into statistical analysis of simulation output data and the simulation user should follow these developments.

ACKNOWLEDGMENT

I would like to express my appreciation to Daniel Wood for his assistance in preparing this paper by his programming and running the simulation model presented in Section IV.

VII. BIBLIOGRAPHY

Adiri, I. and Avi-Itzhak, B., "A Time-Sharing Queue With a Finite Number of Customers," Journal of ACM, Vol. 16, No. 2, April 1969.

Anderson, H.A. and Sargent, R.G., "A Statistical Evaluation of the Scheduler of an Experimental Interactive Computing System," Statistical Computer Performance Evaluation, Academic Press, 1972.

Anderson, H.A. and Sargent, R.G., "Investigation into Scheduling for an Interactive Computing System," IBM Journal of Res. & Dev, Vol. 18, No. 2, March 1974.

Baskett, F., "Confidence Intervals for Simulation Results: A Case Study of Buffer Pool Performance," Proceedings of Computer Science & Statistics: Seventh Annual Symposium on the Interface.

Box, G.E.P. and Jenkins, G.M., Time Series Analysis Forecasting and Control, Holden-Day, 1970.

Conway, R.W., "Some Tactical Problems in Digital Simulation," Management Science, Vol. 10, No. 1, October 1973.

Crane, M.A. and Iglehart, D.L., "Simulating Stable Stochastic Systems, I: General Multiserver Queues," Journal of the ACM, Vol. 21, No. 1, January 1974.

Crane, M.A. and Iglehart, D.L., "Simulating Stable Stochastic Systems, II: Markov Chains," Journal of the ACM, Vol. 21, No. 1, January 1974.

Crane, M.A. and Iglehart, D.L., "Simulating Stable Stochastic Systems, III: Regenerative Processes and Discrete-Events Simulations," Operations Research, Vol. 23, No. 1, January-February 1975.

Crane, M.A. and Iglehart, D.L., "Simulating Stable Stochastic Systems, IV: Approximation Techniques," Management Science, Vol. 21, No. 11, July 1975.

Devor, R.E., Fisher, R.R. and Dessouky, M.F., "Analysis of Simulation Output Characteristics by Autoregressive-Integrated Moving Average Models," Technical Report, Department of Mechanical and Industrial Engineering, University of Illinois at Urbana-Champaign, 1972.

Duket, S.D. and Pritsker, A.A.B., "Spectral Methods for Simulation Output," Technical Paper, Pritsker and Associates.

Emshoff, J.R. and Sisson, R.L., Design and Use of Computer Simulation Models, The MacMillan Company, 1970.

Fishman, G.S., "The Analysis of Simulation-Generated Time Series," Management Science, Vol. 13, No. 7, March 1967.

Fishman, G.S., Concepts and Methods in Discrete Event Digital Simulation, John Wiley and Sons, 1973.

Fishman, G.S., "Batch Means in Digital Simulation," Technical Report No. 75-7, Curriculum in Operations Research and System Analysis, University of North Carolina at Chapel Hill, 1975.

Bibliography: (Cont'd)

Hunt, A.W., "Statistical Evaluation and Verification of Digital Simulation Models Through Spectral Analysis," Ph.D. Dissertation, The University of Texas at Austin, 1970.

Iglehart, D.L., "Simulating Stable Stochastic Systems, V: Comparison of Ratio Estimators," Naval Res. Logist. Quart., Vol. 22, No. 3, September 1975.

Iglehart, D.L., "Simulating Stable Stochastic Systems, VI: Quantile Estimation," Journal of the ACM, Vol. 23, No. 2, April 1976.

Jenkin, G.M and Watts, D.G., Spectral Analysis and Its Application, Holden-Day, 1968.

Kleijnen, J.P.C., Statistical Techniques in Simulation: Part I and Part II, Marcel Dekker, 1974 and 1975.

Lavenberg, S.S. and Slutz, D.R., "Introduction to Regenerative Simulation," IBM Journal of Res. and Dev., Vol, 19, No. 5, September 1975.

Law, A.M., "Confidence Intervals in Discrete Event Simulation: A Comparison of Replication and Batch Means," Technical Report No. 76-13, Department of Industrial Engineering, University of Wisconsin-Madison, 1976.

MISSION
of
Rome Air Development Center

RADC plans and conducts research, exploratory and advanced development programs in command, control, and communications (C^3) activities, and in the C^3 areas of information sciences and intelligence. The principal technical mission areas are communications, electromagnetic guidance and control, surveillance of ground and aerospace objects, intelligence data collection and handling, information system technology, ionospheric propagation, solid state sciences, microwave physics and electronic reliability, maintainability and compatibility.