Proceedings of the 1987 Winter Simulation Conference
A. Thesen, H. Grant, W. David Kelton (eds.)

MULTIVARIATE INFERENCE IN STATIONARY SIMULATION USING BATCH MEANS

Der-Fa R. Chen
Department of Business Management
National Sun-Yat Sen University
Kaosiung, Taiwan

and

Andrew F. Seila
Department of Management Science
and Information Technology
University of Georgia

ABSTRACT

Methods are presented for computing a joint confidence region and simultaneous confidence intervals for the mean of multivariate observations from a simulation operating in steady-state. These methods place observations into batches and use the batch means to compute the confidence region or confidence intervals. As sequential method to determine the batch size that assures independence among the multivariate batch means is presented and empirical studies are used to demonstrate the performance and properties of these methods.

1. INTRODUCTION

Suppose that a simulation operating in steady-state produces a sequence of vector-valued observations X_1, X_2, ..., where each X_i has p components. $X_i = (X_{i1}, X_{i2}, \ldots, X_{ip})$. A joint confidence region or simultaneous confidence intervals are desired for $\mu = E(X_i)$.

For example, consider a multi-item inventory system in which items are substitutable or otherwise interact. Let X_{ij} be the amount of item j in the inventory at the end of period i. Then, μ is the vector of mean inventory levels. This estimation problem also occurs in simulations where, each time an entity exits the system, one can record multiple performance measure observations related to that entity. For example, in a manufacturing system, one might wish to estimate jointly the mean processing time per item and the mean time required for a special part of the processing.

A method of estimating μ is proposed which uses batched means of vectors of observations. Section 2 describes the proposed methodology; section 3 discusses determining an appropriate batch size; and section 4 presents the results of some empirical testing of the method.

2. MULTIVARIATE BATCH MEANS

Various papers (Fishman, 1978; Law, 1977) have discussed batch means for univariate observations. The objective of this paper is to propose an extension of this methodology to multivariate observations. Suppose that one has n vectors of observations X_1, X_2, ..., X_n, and that n = mk, where m > 0 and k > 0 are integers. The entire series then can be divided into m batches of k observations each. The i^{th} batch consists of observations $X_{(i-1)k+1}$, ..., X_{ik}. The i^{th} batch mean, then is:

$$\bar{X}_i = \frac{1}{k} \sum_{j=1}^{k} X_{(i-1)k+j} .$$

Under appropriate conditions, one can show that as $k \to \infty$, $\bar{X}_i$ and $\bar{X}_j$, $i \neq j$, are uncorrelated. If the batch size, k, is assumed to be large enough that this is approximately true, then the m batch means can be treated as if they are uncorrelated vectors of observations, and standard multivariate confidence region estimators can be applied.

In particular, the point estimator for μ is the sample mean of the batch means:

$$\hat{\mu} = \frac{1}{m} \sum_{i=1}^{k} \bar{X}_i .$$

If $\underset{\sim}{S}$ denotes the sample variance-covariance matrix for $\bar{X}_1, \bar{X}_2, \ldots, \bar{X}_m$:

$$S = \frac{1}{m-1} \sum_{i=1}^{m} (X_i - \bar{X}_i)(X_i - \bar{X}_i)' ,$$

where X' denotes the transpose of X. A 100(1-α)-percent confidence region for μ is then given by all vectors x which satisfy:

$$m(\hat{\mu} - x)' \, S^{-1}(\hat{\mu} - x) \leq \frac{p(m-1)}{m-p} F_{\alpha;p,m-p} ,$$

where $F_{\alpha;p,m-p}$ is the 100(1-α)-percentage point of the F-distribution with p and m-p degrees of freedom in the numerator and denominator, respectively. Other methods are available for computing simultaneous confidence intervals for μ; see Morrison (1976).

3. DETERMINING THE BATCH SIZE

This batch means method depends, as does the univariate batch means method, upon establishing a batch size that is large enough that all batch means from distinct batches are approximately uncorrelated. Since the observations X_1, X_2, ... are from a stationary process, the sequence of batch means $\bar{X}_1$, $\bar{X}_2$, ..., $\bar{X}_m$ are from a stationary stochastic process. Let $\Gamma(j)$ denote the autocovariance function of the batch means process:

$$\Gamma(j) = E \{ \bar{X}_i \bar{X}'_{i+j} \} .$$

The batch size , k, must then be made large enough that $\Gamma(j) = 0$, for j=1, 2, ...

Fishman (1979) has presented a method for determining the batch size for univariate processes. The multivariate generalization of this method is complicated somewhat by the fact that $\Gamma(j)$ includes cross-correlations among components of $\bar{X}_i$ as well as lag-j autocorrelations. The approach used in this research was to assume that the batch mean process is a multivariate autoregressive process of order p Hannan, (1979):

$$\bar{X}_i + B_1\bar{X}_{i-1} + \dots + B_p\bar{X}_{i-p} = \epsilon_i ,$$

where $B_1, B_2, \dots, B_p$ are p x p matrices and ϵ_i are i.i.d. random vectors with mean 0 and covariance matrix Σ_ϵ.

For a multivariate batch means process, two assumptions seem reasonable to make: First, if the batches are large enough that adjacent batch means are uncorrelated, then they are large enough that a central limit theorem applies, and $\bar{X}_i$ are approximately normally distributed. Secondly, the lag-1 autocovariance dominates the others. Therefore, it is reasonable to fit the batch means process to a first-order autoregressive model:

$$\bar{X}_i - B_1\bar{X}_{i-1} = \epsilon_i , \ i=1,2,3, \dots , m.$$

If this is the case, then B_1 can be estimated by solving:

$$\hat{B}_1 = C_1C_o^{-1} ,$$

where C_j is the estimate of the autocovariance at lag j.
Let:

$$\Omega(k,p) = \frac{|C_o - \hat{B}_1{}'C_o\hat{B}_1|}{|C_o|} ,$$

where $|A|$ denotes the determinant of matrix A. Then, under the null hypothesis that $\bar{X}_1, \bar{X}_2, \dots, \bar{X}_m$ are serially independent (i.e., $B_1 = 0$) the test statistic $-(k-p)\log \Omega(k,p)$ is approximately distributied as a chi-square random variable with p^2 degrees of freedom. Thus, if $-(k-p)\log \Omega(k,p) > \chi^2_{p^2}(\alpha)$, where $\chi^2_{p^2}(\alpha)$ is the upper $(1-\alpha)$-percentage point of the chi-square distribution, then the conclusion will be drawn that the batch means are not uncorrelated, and the batch size is not sufficiently large.

To determine a appropriate batch size, the number of observations per batch was initially set to 20 and the test for serial independence was applied. If the hypothesis was rejected, the batch size was doubled and the test reapplied. This process continued until either the null hypothesis was accepted or the number of batches was less than 8. This sequential method is the multivariate analog of the approach described in Fishman [1].

4. EMPIRICAL STUDIES

This methodology was tested on the University of Georgia's CYBER 170/750 computer. The system that was simulated was a random queue with two servers, infinite capacity in each queue, Poisson arrivals and independent exponential service times. The parameters that were estimated were: μ_1, the expected waiting time in queue 1; μ_2, the expected waiting time in queue 2; and μ_3, the expected total time in system. Clearly, waiting times in the first queue are independent of waiting times in the second queue; however, waiting times in each queue are correlated with total time in system.

Initial sample sizes were set at 2000, 3000, 4000 and 5000 observations, and the simulation was run with traffic intensities at stage 1 of 0.9 (high), 0.75 (moderate) and 0.5 (low). For each case, 100 replications were run. Roy-Bose confidence regions were computed using confidence coefficient .975. The results are shown in Tables 1 and 2. The failure rate in Table 1 is the proportion of replications in which the test for uncorrelated batch means failed to accept the null hypothesis and conclude that the batch means are uncorrelated. The conditional coverage rate is the proportion of confidence regions that covered the true parameter vector, omitting those replications where the batch size test failed. The overall coverage rate is the proportion of replications, out of 100, in which both an appropriate batch size was found and the true parameter vector was covered. Since the number of replications was 100, these proportions can be assumed to have an error of approximately ± .08.

These simulations show two things: First, as one would expect, much larger sample sizes are required to compute reliable confidence regions for more congested systems. This is understandable since the observations are more highly correlated and thus each observation carries less information. Secondly, required sample sizes for multivariate estimation are much larger than for univariate estimation. With sufficiently large sample sizes, however, this approach does produce reliable confidence regions. Additional simulation studies are being done to compare alternative methods for computing confidence regions and simultaneous confidence intervals.

REFERENCES

Fishman, G. S. (1978). Grouping Observations in Digital Experiments. Management Science 24, 510-521.

Fishman, G. S. (1979). Principles of Discrete Event Simulation. John Wiley, New York.

Hannan, E. J. (1970). Multiple Time Series. John Wiley, New York.

Law, A. M. (1977). Confidence Intervals in Discrete Event Simulation: A Comparison of Replication and Batch Means. Naval Research Logistics Quarterly 24, 667-678.

Morrison, D. F. (1976). Multivariate Statistical Methods. McGraw-Hill, New York.

Table 1: Simulation Results for High Traffic Intensity

Sample Size	Overall Coverage Rate	Failure Rate	Conditional Coverage Rate
2000	.62	.11	.70
3000	.67	.08	.73
4000	.72	.03	.74
5000	.73	.05	.77

Table 2: Simulation Results for Moderate and Low Traffic Intensity

Sample Size	Coverage Rates Moderate	Low
2000	.80	.94
3000	.89	.95
4000	.87	.95
5000	.91	.97

AUTHORS' BIOGRAPHIES

Der-Fa R. Chen is Assistant Professor of Business Management at National Sun-Yat Sen University in Taiwan. He completed the MBA and Ph.D. degree at The University of Georgia in 1981 and 1985, respectively. His professional interests are in the areas of simulation methodology and stochastic models.

Der-Fa Robert Chen
Department of Business Management
National Sun Yat-Sen University
Kaohsiung, Taiwan 800
Republic of China
(07) 531-6171

Andrew F. Seila is an Associate Professor of Management Sciences in the College of Business Administration at the University of Georgia. He received the B.S. degree in physics in 1970 and the Ph.D. degree in operations research in 1976, both from The University of North Carolina at Chapel Hill. Before joining the faculty at the University of Georgia, he was at Bell Laboratories in Holmdel, New Jersey. His current research interests include all aspects of discrete event simulation, including model development and validation, and simulation output analysis, especially analysis of multivariate output data. He is a member of ORSA, TIMS and ASA.

Andrew F. Seila
Department of Management Sciences
and Information Technology
297 Brooks Hall
University of Georgia
Athens, Georgia 30602 USA
(404) 542-3586