Proceedings of the 1996 Winter Simulation Conference
ed. J. M. Charnes, D. J. Morrice, D. T. Brunner, and J. J. Swain

USING NULL EFFECTS ESTIMATORS IN THE ANALYSIS OF SIMULATION OUTPUT

Lee W. Schruben
School of OR&IE
Cornell University
Ithaca, NY 14853

David Goldsman
School of ISyE
Georgia Institute of Technology
Atlanta, GA 30332

Douglas J. Morrice
MSIS Department
The University of Texas at Austin
Austin, Texas 78712-1175
End to End Simulation Department
Schlumberger Austin Research
Austin, Texas 78720

ABSTRACT

In a simulation experiment there are certain *null effects* whose true values are known to be zero. This paper investigates various estimators of the null effects. We apply these estimators of zero in two ways. First, null effects estimators are used to construct confidence intervals. We derive a null effects version of the batch means confidence interval that has certain desirable properties. Second, null effects estimators are used as control variates in variance reduction schemes. The achieved variance reductions are modest, but we present some interesting applications; for example, the output from one simulation model can be used to compute control variates for the output of another.

1 INTRODUCTION

This paper deals with various ways that simulation output can be used to estimate *null effects*, i.e., quantities whose true values are known to be zero. We subsequently apply these estimators of zero in two ways.

First, we will use the null effects to develop a confidence interval estimator (c.i.e.) for the mean of a stationary simulation output process. This c.i.e. is similar to the familiar batch means (BM) c.i.e., but it is obtained via a different route. (Indeed, this paper is not intended to change the way in which simulation experiments are run and analyzed but rather to present an alternative way of thinking about simulation output analysis!) The null effects interval estimation procedure does have desirable properties, one of which should facilitate in the selection of a batch size.

Next we show how null effects estimators can be used as control variates (c.v.'s) in variance reduction schemes. In one scheme, null effects estimators serve as internal c.v.'s and yield estimators of the mean that have smaller variance than the sample mean. A second application uses the output from a "control" simulation to compute a c.v. for an estimator from another simulation. The true mean of the control system need not be known. A third variance reduction scheme indirectly uses null effects c.v.'s to improve the estimator of the process mean.

2 BACKGROUND AND NOTATION

This section gives some preliminary definitions and examples. We shall regard the stationary stochastic sequence $Y_1, \ldots, Y_n$ generated by running a discrete-event simulation program as a vector $\boldsymbol{Y} \equiv (Y_1, \ldots, Y_n)'$ in R^n. This paper demonstrates the potential advantages of viewing the simulation output as consisting of different components of orthogonal basic vectors. In order to facilitate the study of various estimators of the variability of the simulation output, $\boldsymbol{Y}$ will be projected onto different coordinate systems that are determined by sets of basic vectors.

Sometimes a basis may be intentionally chosen to be *incomplete*; i.e., the basis only spans a proper subspace R^b of R^n $(b < n)$. We denote a basis for R^b as b basic vectors, $\boldsymbol{\omega}_i \equiv (\omega_{i,1}, \ldots, \omega_{i,n})'$, $i = 1, \ldots, b$, each in R^n. The i-th component of $\boldsymbol{Y}$ in this basis is the orthogonal projection of $\boldsymbol{Y}$ and is given by

$$Y_i^\omega \equiv \frac{\boldsymbol{\omega}_i'\boldsymbol{Y}}{\boldsymbol{\omega}_i'\boldsymbol{\omega}_i} = \frac{\sum_{j=1}^n \omega_{i,j} Y_j}{\sum_{j=1}^n \omega_{i,j}^2}.$$

If $E[Y_i^\omega] = 0$, we call Y_i^ω a *null effect*. We will see that choosing a basis with components that are null effects (referred to here as a *null effects basis*) has certain advantages.

2.1 Example 1: Batch Means

BM is a popular method of estimating c.i.e.'s for $\mu \equiv E[Y_j]$. The BM basis $\boldsymbol{\omega}_i$, $i = 1, \ldots, b$, is not a null effects basis; its components are

$$\omega_{i,j} = \begin{cases} 1 & \text{if } (i-1)m+1 \le j \le im \\ 0 & \text{otherwise,} \end{cases}$$

with $m = n/b$ (assume b divides n). This basis is incomplete if $b < n$. Further, the subspace spanned by the basis changes radically when b changes (with fixed n); see Section 3.1.

2.2 Example 2: Harmonic Analysis

In order to identify specific periodic oscillations in the output vector, $\boldsymbol{Y}$ can be projected onto the Fourier basis (Bloomfield, 1976). The basic vectors are (assuming n is even):

$$\omega_{i,j} = \begin{cases} (2/n)\cos(2\pi ij/n), & i = 1,\ldots,\frac{n}{2} \\ (2/n)\sin(2\pi(i-\frac{n}{2})j/n), & i = \frac{n}{2}+1,\ldots,n. \end{cases}$$

Since $\boldsymbol{\omega}_i'\boldsymbol{\omega}_k = 0$ for $i \neq k$ and $\sum_j \omega_{i,j} = 0$ for all i (Chatfield 1980, p. 130), the Fourier basis is orthogonal and $E[Y_i^\omega] = 0$ for all i for a stationary sequence $\boldsymbol{Y}$. So the Fourier basis is a complete null basis.

2.3 Example 3: The Walsh Basis

The vectors in the Walsh basis are the columns in the design matrix for a full factorial experimental design with $n = 2^k$, including all additive effects and interactions (Beauchamp 1975, Sanchez 1986). For example, if a simulation run produces 8 observations, the (complete) Walsh basis is

$$\begin{aligned}
\boldsymbol{\omega}_1 &= (1,1,1,1,1,1,1,1)' \\
\boldsymbol{\omega}_2 &= (1,1,1,1,-1,-1,-1,-1)' \\
\boldsymbol{\omega}_3 &= (1,1,-1,-1,-1,-1,1,1)' \\
\boldsymbol{\omega}_4 &= (1,1,-1,-1,1,1,-1,-1)' \\
\boldsymbol{\omega}_5 &= (1,-1,-1,1,1,-1,-1,1)' \\
\boldsymbol{\omega}_6 &= (1,-1,-1,1,-1,1,1,-1)' \\
\boldsymbol{\omega}_7 &= (1,-1,1,-1,-1,1,-1,1)' \\
\boldsymbol{\omega}_8 &= (1,-1,1,-1,1,-1,1,-1)'
\end{aligned}$$

The variable $\boldsymbol{\omega}_1$ corresponds to the mean effect; $\boldsymbol{\omega}_2$, $\boldsymbol{\omega}_4$, and $\boldsymbol{\omega}_8$ correspond to the main effects in a saturated $n = 2^3$ factorial design with an additive model. The ordering of the Walsh basis is by *sequency* (i.e., $\boldsymbol{\omega}_i$ has $i-1$ sign changes). For stationary $\boldsymbol{Y}$, the Walsh basis is orthogonal and $E[W_i] = 0$, $i = 2,\ldots,b$, where $W_i \equiv Y_i^\omega$ for all i. So the Walsh basis is a null effects basis.

3 APPLICATIONS

3.1 Confidence Interval Estimation

Here we discuss two c.i.e.'s for μ based on batch means and Walsh components.

3.1.1 Batch Means c.i.e.'s

The BM $100(1-\alpha)\%$ c.i.e. for μ is

$$\mu \in \bar{Y}_n \pm t_{\alpha/2,b-1}\sqrt{\frac{\sum_{i=1}^b(\bar{Y}_{i,m}-\bar{Y}_n)^2}{(b-1)b}}, \quad (1)$$

where the i-th batch mean $\bar{Y}_{i,m} \equiv Y_i^\omega = \sum_{j=(i-1)m+1}^{im} Y_j/m$, $i = 1,\ldots,b$, the sample mean $\bar{Y}_n \equiv \sum_{i=1}^n Y_i/n = \sum_{i=1}^b Y_i^\omega/b$, and $t_{\gamma,\nu}$ is the upper-γ quantile of the t-distribution with ν degrees of freedom.

3.1.2 Walsh c.i.e.'s

A c.i.e. for μ based on the Walsh components is

$$\mu \in W_1 \pm t_{\alpha/2,b-1}\sqrt{\frac{\sum_{i=2}^b W_i^2}{b-1}}. \quad (2)$$

The asymptotic validity of this c.i.e. is discussed next.

3.1.3 Asymptotic Justification

The Walsh components, W_i, $i = 1,\ldots,b$, can be formed from non-overlapping batches of $\boldsymbol{Y}$. For instance, if $b = 2$, then $W_1 = (\bar{Y}_{1,n/2}+\bar{Y}_{2,n/2})/2$ and $W_2 = (\bar{Y}_{1,n/2}-\bar{Y}_{2,n/2})/2$. Further, when *$b$ is a power of 2 and evenly divides n*, the BM and Walsh c.i.e.'s, (1) and (2), are *equal*. We prove this "n-dimensional Pythagorean theorem" for the special case of $b = n$ ($b \le n$ is similar, but tedious): Observe that for the Walsh basis,

$$\sum_{i=1}^n \omega_{i,j}\omega_{i,k} = \sum_{i=1}^n \omega_{j,i}\omega_{k,i} = \boldsymbol{\omega}_j'\boldsymbol{\omega}_k = \begin{cases} n, & j=k \\ 0, & j \neq k. \end{cases}$$

Then

$$\sum_{i=1}^n W_i^2 = \frac{1}{n^2}\sum_{i=1}^n \boldsymbol{Y}'\boldsymbol{\omega}_i\boldsymbol{\omega}_i'\boldsymbol{Y} = \frac{1}{n}\sum_{i=1}^n Y_i^2.$$

So

$$n\sum_{i=2}^n W_i^2 = \sum_{i=1}^n Y_i^2 - nW_1^2 = \sum_{i=1}^n (Y_i - \bar{Y}_n)^2,$$

and thus, (1) equals (2). □

Consider $\boldsymbol{Y}$ as a scaled continuous-time function by defining $Y(t) \equiv Y_{\lfloor nt \rfloor}$, $0 \le t \le 1$, where $\lfloor \cdot \rfloor$ is the greatest integer function. The projection basis now consists of b continuous-time functions, $\{\omega_i(t), 0 \le t \le 1\}$, $i = 1, \ldots, b$, with $\omega_i(t) \equiv \omega_{i,\lfloor nt \rfloor}$. Then

$$Y_i^\omega \equiv \frac{\boldsymbol{\omega}_i' \boldsymbol{Y}}{\boldsymbol{\omega}_i' \boldsymbol{\omega}_i} = \frac{\int_0^1 \omega_i(t) Y(t)\, dt}{\int_0^1 \omega_i^2(t)\, dt}.$$

We apply the weak convergence arguments of Schruben (1983) and Glynn and Iglehart (1990) to obtain the asymptotic i.i.d. normality of the W_i's: With b fixed, the Walsh basis functions, $\{\omega_i(t), 0 \le t \le 1\}$, are constant over non-overlapping intervals. If the smallest such interval is selected as a batch size, then under general conditions, the resulting batch means converge weakly to i.i.d. normality; i.e., $\sqrt{m}(\bar{Y}_{i,m} - \mu) \xrightarrow{\mathcal{D}} \text{Nor}(0, \sigma^2)$, $i = 1, \ldots, b$, where $\sigma^2 \equiv \lim_{n \to \infty} n \text{Var}(\bar{Y}_n)$ and $\xrightarrow{\mathcal{D}}$ denotes convergence in distribution as $m \to \infty$. Since the Walsh basis components are orthogonal linear combinations of asymptotically i.i.d. normal batch means, they are also asymptotically i.i.d. normal; i.e., $\sqrt{m} W_i \xrightarrow{\mathcal{D}} \text{Nor}(0, \sigma^2/b)$, $i = 2, \ldots, b$.

3.1.4 Practical Considerations

Recall that when b is a power of two and evenly divides n, (1) and (2) are equal. In fact, based on empirical experimentation (not reported here), the performance characteristics of the BM and Walsh c.i.e.'s are essentially the same in general.

As is well known by simulation users, the main difficulty in using BM lies in determining the proper b to use (with n fixed) so that the assumption of i.i.d. normal batch means is approximately met. Schmeiser (1982) effectively argues that we rarely need $b > 30$. Sequential algorithms for finding an appropriate b (e.g., Fishman 1978, Fishman and Yarberry 1995, and Law and Kelton 1982) usually start with a particular value of b and reduce it while performing statistical tests of correlation and/or normality on the batch means.

A possible advantage of the Walsh c.i.e. procedure is that the Walsh basis is *constructive* as b increases; i.e., previous Walsh basic vectors remain in the basis as more are added. This contrasts with the BM basis which can change radically as b changes (with n fixed). Perhaps this observation could be used as the foundation for a batch size algorithm based on Walsh components.

We also see that the Walsh basis uses all of the output to estimate each W_i; the W_i's may converge to normality more quickly than do batch means. Indeed, Q-Q plots against normality have indicated consistently better performance for the W_i's (in terms of normality) than for the $\bar{Y}_{i,m}$'s.

3.2 Variance Reduction

Since the null effects have known means of zero, they make natural c.v.'s to use in variance reduction schemes for simulation. We give some applications.

3.2.1 Null Effects c.v.'s in a Single Run

When there is positive serial correlation in the stationary output series $\boldsymbol{Y}$, it is possible to construct a linear estimator of $\mu = E[Y_j]$ that is superior to the sample mean W_1. Several such estimators have been proposed (Halfin 1982) but usually in rather restrictive situations that require knowledge of the covariance structure of $\boldsymbol{Y}$. Our estimators require little knowledge of this structure. We can use c.v.'s of the form $C \equiv W_1 + kW_i$, $i = 2, \ldots, b$. Of course, $E[C] = E[W_1] = \mu$. The constant k is chosen in the hope of minimizing $\text{Var}(C)$. For $k > 0$, $\text{Var}(C) < \text{Var}(W_1)$ iff $\text{Cov}(W_1, W_i) < -k\text{Var}(W_i)/2$.

To illustrate, let $n = 4$ and $C = W_1 + kW_3$. Then $W_1 = (Y_1 + Y_2 + Y_3 + Y_4)/4$ and $W_3 = (Y_1 - Y_2 - Y_3 + Y_4)/4$. If we define $\gamma_j \equiv \text{Cov}(Y_i, Y_{i+j})$ for all j, some easy algebra gives $\text{Cov}(W_1, W_3) = (\gamma_3 - \gamma_1)/8$. For many stochastic processes, γ_j is decreasing in j; then $\text{Cov}(W_1, W_3) < 0$, and so small enough k will yield $\text{Var}(C) < \text{Var}(W_1)$. Now consider a first-order autoregressive [AR(1)] process, $Y_i = \mu + \alpha(Y_{i-1} - \mu) + \epsilon_i$, $i = 1, 2, \ldots$, with i.i.d. $\text{Nor}(0, 1-\alpha^2)$ ϵ_i's, and initialized in steady state. For the AR(1), $\gamma_j = \alpha^{|j|}$ for all j. We find $\text{Var}(W_1) = (1+\alpha)(2+\alpha+\alpha^2)/8$ and $\text{Var}(C) = (1+\alpha)[2+\alpha+\alpha^2 - 2k\alpha(1-\alpha) + k^2(2-\alpha)(1-\alpha)]/8$. $\text{Var}(C)$ is minimized with respect to k when $k = \alpha/(2-\alpha)$, in which case $\text{Var}(C) = (1/2)(1+\alpha)/(2-\alpha)$ and $\text{Var}(C)/\text{Var}(W_1)$ is between 0.964 and 1 for $\alpha \in [0, 1]$; so the variance reduction is modest.

3.2.2 Null Effects External c.v.'s

In external c.v. applications where one (stationary) simulation is used to produce c.v.'s for a second simulation, the expected value of the control simulation's variate must be known. Null effects estimators can therefore be used as c.v.'s when little is known about either simulation. The effectiveness of the approach depends on there being positive serial correlation within each of the individual output series (which is typical of many simulated systems) and positive correlation between outputs from the two simulations (as can often be induced by using common pseudorandom number streams as input).

3.2.3 Indirect Use of Null Effects c.v.'s

Here we indirectly use null effects c.v.'s to improve the estimator of the mean. The rationale for selecting the form of the c.v. used is as follows. Suppose a "reasonable" hypothesis involving μ is H_O versus H_A, and we know that H_O is true for the time series $\boldsymbol{Y}$ that is used to estimate μ. After estimating μ, we *alter* the time series by computing a null effect (which makes H_A true). The test statistic for the hypothesis test should be highly correlated with an indicator variable of whether H_O or H_A is true (else the test would not be very powerful). We should therefore be able to use the test statistic computed from the altered time series as a c.v. for the estimator from the original series.

For instance, consider the AR(1) process with $n = 4$. Suppose the altered output is $-\boldsymbol{\omega}_2'\boldsymbol{Y} = (-Y_1, -Y_2, Y_3, Y_4)'$. This altered series of four observations could be considered to contain severe initialization bias. A reasonably powerful test for initialization bias is based on the area under the so-called *standardized time series* (Schruben, Singh, and Tierney 1983). This test statistic is $\widetilde{Y} \equiv (3Y_1 + Y_2 + Y_3 + 3Y_4)/2$, for which $E[\widetilde{Y}] = 4\mu$. Thus, an unbiased c.v. estimator for μ is $C \equiv (\widetilde{Y} - \bar{Y}_n)/3$. Some algebra shows that $\mathrm{Var}(C) < \mathrm{Var}(\bar{Y}_n)$ iff $\gamma_0 - \gamma_2 < 2(\gamma_1 - \gamma_3)$ iff $\alpha > 1/2$. However, the variance reduction from this example is again modest. For example, when $\alpha = 0.9$, we obtain a variance reduction of about 4%. This is consistent with the variance reductions found by Halfin when the covariance structure is known and an optimal linear estimator is used.

4 CONCLUDING REMARKS

We have demonstrated that viewing a simulated time series as a vector in n-dimensional space can present useful insights when different bases for the space are considered. Relationships between the BM and Walsh bases yield alternative procedures for estimation and variance reduction.

Some subjects for future research: Techniques for consistent estimation of the Walsh spectrum (perhaps similar to spectral windowing) could be applied to BM estimation, and new procedures for selecting a batch size might be engineered using the Walsh basis. Also, although we illustrated some applications of null effects c.v.'s, no attempt were made to generalize the results, optimize the control, or use multiple c.v.'s — all ideas that might improve the results.

Certainly, the use of a different basis for the simulation output vector is one way of viewing the data differently, and it can lead to some entertaining new ideas.

ACKNOWLEDGEMENTS

This research was supported by National Science Foundation Grant No. DMI-9622269. The third author gratefully acknowledges the financial support of the CBA/GSB Faculty Research Committee of the College of Business of the University of Texas at Austin and the End to End Simulation Department at Schlumberger Austin Research.

REFERENCES

Beauchamp, K. G. 1975. *Walsh Functions and Their Applications.* New York: Academic Press.

Bloomfield, P. 1976. *Fourier Analysis of Time Series: An Introduction.* New York: John Wiley and Sons.

Chatfield, C. 1980. *The Analysis of Time Series: An Introduction.* New York: Chapman and Hall.

Fishman, G. S. 1978. *Principles of Discrete Event Simulation.* New York: John Wiley and Sons.

Fishman, G. S. and L. S. Yarberry. 1995. An Implementation of the Batch Means Method. Technical Report, Department of Operations Research, University of North Carolina, Chapel Hill, NC.

Glynn, P. W. and D. L. Iglehart. 1990. Simulation Output Analysis Using Standardized Time Series. *Mathematics of Operations Research* 15:1–16.

Halfin, S. 1982. Linear Estimators for a Class of Stationary Queueing Processes. *Operations Research* 30:515–529.

Law, A. M. and W. D. Kelton. 1982. Confidence Intervals for Steady-State Simulation, II: A Survey of Sequential Procedures. *Management Science* 28:550–562.

Sanchez, P. J. 1986. On the Correspondence Between Walsh Spectral Analysis and 2^k Factorial Experimental Designs. In *Proceedings of the 1986 Winter Simulation Conference* (ed. J. R. Wilson, J. O. Henriksen, and S. D. Roberts), 306–312. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey.

Schmeiser, B. W. 1982. Batch Size Effects in the Analysis of Simulation Output. *Operations Research* 30:556–568.

Schruben, L. W. 1983. Confidence Interval Estimation Using Standardized Time Series. *Operations Research* 31:1090–1108.

Schruben L. W., H. Singh, and L. Tierney. 1983. Optimal Tests for Initialization Bias in Simulation Output. *Operations Research* 31:1167–1178.

AUTHOR BIOGRAPHIES

LEE W. SCHRUBEN is a Professor in the School of Operations Research and Industrial Engineering at Cornell University. He received his undergraduate degree in engineering from Cornell University and his Ph.D. from Yale University. His research interests are in the statistical design and analysis of simulation experiments and in graphical simulation modeling methods.

DAVID GOLDSMAN is an Associate Professor in the School of Industrial and Systems Engineering at the Georgia Institute of Technology. His research interests include simulation output analysis and ranking and selection. He is the Simulation Department Co-Editor for *IIE Transactions*, and the Past-President of the INFORMS College on Simulation.

DOUGLAS J. MORRICE is an Associate Professor in the Department of Management Science and Information Systems at The University of Texas at Austin. He is currently on sabbatical leave in the End to End Simulation Department at Schlumberger Austin Research. Dr. Morrice received his undergraduate degree in Operations Research at Carleton University in Ottawa, Canada. He holds an M.S. and a Ph.D. in Operations Research and Industrial Engineering from Cornell University. His research interests include discrete-event and qualitative simulation modeling and the statistical design and analysis of large-scale simulation experiments. Dr. Morrice is a member of INFORMS and the Council of Logistics Management. He served as Secretary for the INFORMS College on Simulation (1994-1996) and is Co-Editor of the *Proceedings of the 1996 Winter Simulation Conference.*