Proceedings of the 1992 Winter Simulation Conference
ed. J. J. Swain, D. Goldsman, R. C. Crain, and J. R. Wilson

FORCE ALLOCATION THROUGH CONSTRAINED OPTIMIZATION OF STOCHASTIC RESPONSE SURFACES

R. Garrison Harvey

HQ Air Mobility Command,
Command Analysis Group
Scott AFB, Il 62225

Kenneth W. Bauer Jr.

Dept of Operational Sciences
School of Engineering
AF Institute of Technology,
WPAFB, Oh 45433

Joseph R. Litko

HQ Air Mobility Command,
Command Analysis Group
Scott AFB, Il 62225

ABSTRACT

This paper presents an alternative to the traditional approaches to optimizing a stochastic response surface subject to constraints. It focuses on the stochastic nature of the response surface and the implications for the subsequent optimization. This research presents a three step process to evaluate stochastic response surfaces subject to constraints. Step 1 uses experimental design to estimate the response surface and covariance matrix. Step 2 samples the objective function of the linear program (i.e., response surface) and identifies the associated extreme points. Step 3 presents a method to identify the optimal extreme point and present that information to a decision maker.

1 INTRODUCTION

To focus the discussion, consider the hypothetical problem of allocating a fixed budget to purchasing crews for a mixed fleet of aircraft such as Air Mobility Command (AMC) heavy transports -- the C-5, C-17, and C-141. Each type of aircraft has different characteristics (e.g., material handling requirements, crew size, parking space, speed, payload and fuel). Because the planes are reliable, and because of safety limits on the hours a crew may fly, there are actually three or more crews per plane. The exact number of crews per plane (called the crew ratio) which maximizes cargo throughput for a given scenario can be analyzed with a mix of simulation and linear programming.

First, a simulation model would be used to derive a response surface of throughput as a function of crew ratio and other variables. The simulation captures some of the unpredictable events such as maintenance, weather delays, variable service times and varying system capability. Second, an LP model would maximize throughput considering the relevant constraints, e.g., budget for crews, physical constraints at en route stops, etc. Naive analysis treats the response surface as a deterministic function, missing the fact that the coefficients are random variables derived from the simulation model.

The random, or stochastic, nature of the coefficients in the response surface is at the core of this paper. This paper is organized as follows. First, we present some background to the problem. Next, we discuss the impact of the stochastic nature of the response surface in a constrained optimization problem. Then we develop a practical means to identify the "true" optimum point. Finally, we conclude the paper with a summary and offer some recommendations.

2 BACKGROUND

Biles and Swain (1977) present several strategies for constrained simulation optimization. They fit and validate a response surface using an n-dimensional simplex, biradial, or equiradial design. They account for the variance of the error term, but assume the "response surfaces are the expected values of the observed responses." They do not directly account for the stochastic nature of the response surface, but use an iterative method by applying an optimization procedure and then returning to the simulation model until stopping criteria are met. Their procedures include direct search techniques, first-order and second-order response surface procedures

Myers (1989) concludes that, "Many users of RSM allow conclusions to be drawn concerning the nature of a response surface and the location of optimal response without taking into account the distributional properties of the estimated attributes of the underlying response surface."

Morben (1987), in solving a "real world" problem, demonstrates a case where using the expected value of a stochastic objective function leads to an answer that falls outside a 95% confidence bound found through a Monte Carlo analysis. This case clearly demonstrates there is a risk in some situations of using only the expected value, and it makes the case for incorporating some form of stochastic analysis.

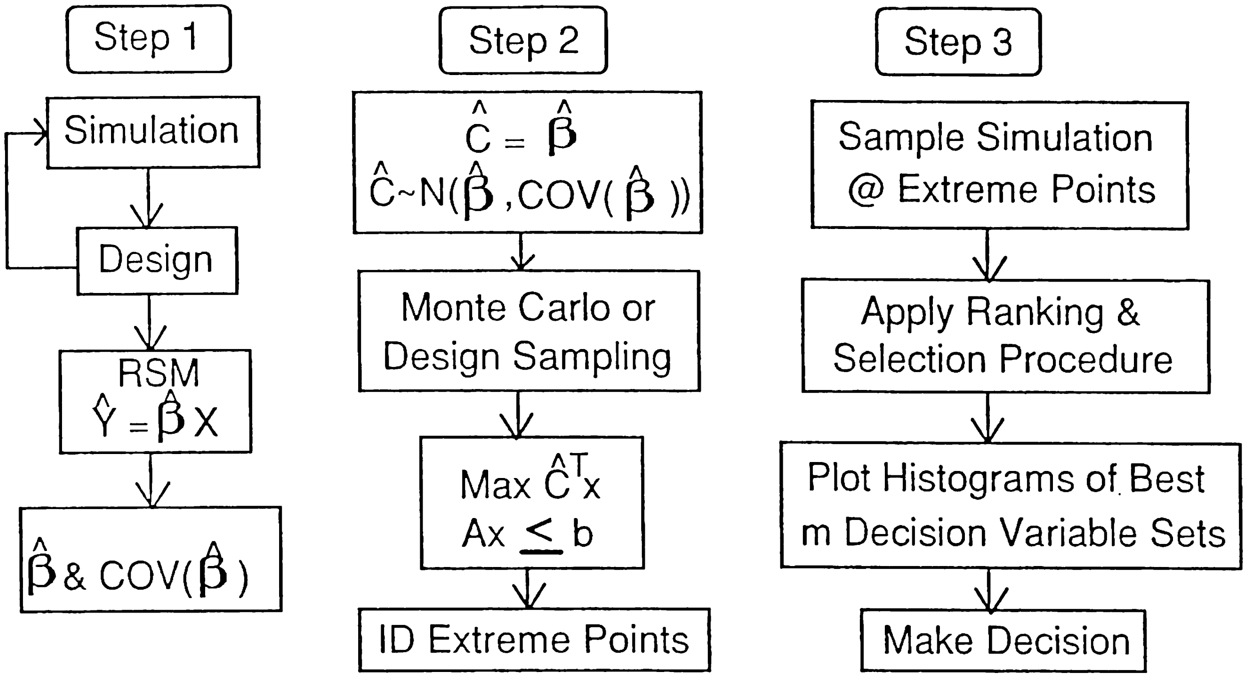

Figure 1: Three Step Approach

3 IMPACT OF ESTIMATION ERRORS

Figure 1 shows the overall layout of our three step approach to optimizing a stochastic response surface subject to constraints. Step 1 is similar to the traditional approach which would estimate β in much the same way. Here, however, we depart. The traditional approach would take the estimates $\hat{\beta}$ and proceed to solve a single linear programming problem whose objective function coefficients are $\hat{\beta}$. In this section we detail the problems which arise from this approach to the problem.

We viewed the simulation as a black box that consists of a "Truth Model" plus noise. Simulation output from a designed experiment allows us to estimate a response surface that becomes the objective function of a linear program. The functions:

$$z^* = LP(\mathbf{c}, A, \mathbf{b}) \quad (1)$$
$$\hat{z}^* = LP(\hat{\mathbf{c}}, A, \mathbf{b}) \quad (2)$$

define the optimal value z^* (or estimated optimal value $\hat{z}^*$) of a linear program. Where

$z^* = \mathbf{c}^T\mathbf{x}$ true (or known) objective function
$\hat{z}^* = \hat{\mathbf{c}}^T\mathbf{x}$ estimated objective function
$\mathbf{x} = (x_1, x_2, \ldots, x_p)^T$ =vector of decision variables
A = constraint matrix
$\mathbf{b}$ = right hand side vector.
$\mathbf{c}$ = true surface coefficients underlying the metamodel
$\hat{\mathbf{c}} = \mathbf{c} + \varepsilon$ estimated coefficients of objective function (response surface), with $\varepsilon \sim N(0, \sigma^2(X^TX)^{-1})$

where X is the design matrix used to estimate the response surface. Since we assume there is no bias in the estimation of $\hat{\mathbf{c}}$, we can obtain z^* (the "true" optimum) if we run the LP with the expected value of the objective function coefficients:

$$z^* = LP(E(\hat{\mathbf{c}}, A, \mathbf{b})) \quad (3)$$

In general, z^* is not equal to the expected value of the linear program with respect to the objective function coefficients. That is,

$$z^* = LP(\mathbf{c}, A, \mathbf{b}) = LP(E(\hat{\mathbf{c}}, A, \mathbf{b})) \neq E(LP(\hat{\mathbf{c}}, A, \mathbf{b})) = E(\hat{z}^*) \quad (4)$$

Further, as the standard error in the estimates of the coefficients increases, the bias in $\hat{z}^*$ and the variance of $\hat{z}^*$, $\sigma^2(\hat{z}^*)$, increases.

The objective function of a linear programming maximization problem is a piecewise linear convex function of the objective function coefficients, $\mathbf{c}$. Looking at this in one dimension for simplicity, as in Figure 2, consider z as a function of the coefficient c_i.

The slope of each piecewise linear segment, α_k, is simply the value of the decision variable x_i in the basis that applies in region k. As the curve indicates, decreasing values of c_i are accompanied by decreasing slopes, i.e. the variable x_i is basic at a lower value, if it is basic at all. If c_i is small enough the ith variable becomes nonbasic and further reductions in c_i have no effect. So too, increasing c_i eventually loses its effect when x_i can be increased no further.

Assume the true objective function value of c_i, call it c_{ik}, lies in the k^{th} piecewise linear interval. Then,

$$\text{for } c_i \geq c_{ik} \quad \text{we have } z(c_i) \geq \alpha_k c_i + z(0)$$
$$\text{for } c_i \leq c_{ik} \quad \text{we have } z(c_i) \geq \alpha_k c_i + z(0)$$

due to the convexity of z with respect to c_i. Therefore, $E(z(c_i)) \geq E(\alpha_k c_i + z(0)) = z(E(c_i)) = z(c_{ik})$, illustrates the point made in equation 4. Figure 2 also illustrates the point that the bias, which is proportional to the shaded area, will also be proportional to the variance in the coefficients. A tight distribution will seldom produce estimates that cross into adjacent piecewise linear segments. The opposite is true of wildly varying estimates of the objective function coefficients.

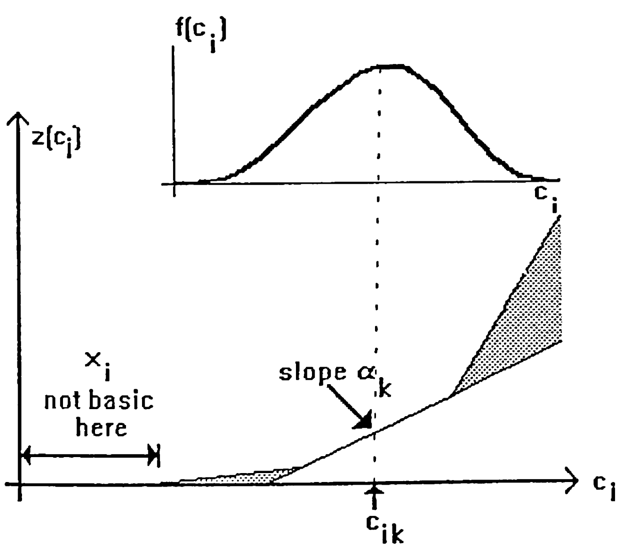

Figure 2: Noise Impact & Bias

With these observations in mind, we performed experiments to illustrate the bias identified in Equation 4. The computer program samples from a "Truth Model" with noise, generates a response surface, and then uses it as the objective function of the linear program. Solving the LP for the estimated optimal value and the estimated optimal extreme point, these are then compared to the true optimal extreme point and its objective function value. Varying the amount of noise added to the problem illustrates its effect on the solutions. We investigated many different types of problems. A variety of problems are described in Harvey (1992). The results presented here are from a linear program with four variables and three constraints, where the coefficients of the objective function are 15, 17, 18, and 20.

Either positive or zero bias in the estimate of the mean was present in all linear programming problems analyzed as in Harvey (1992). Also, as the noise level increased in a given problem the bias increased in a roughly linear trend; here the bias is the mean estimated optimum minus the true optimum.

$$\text{Bias} = E(\hat{z}^*) - z^* \qquad (5)$$

Figures 3 and 4 illustrate a typical case where the standard error = σ(objective function coefficients), the sample size is 1000.

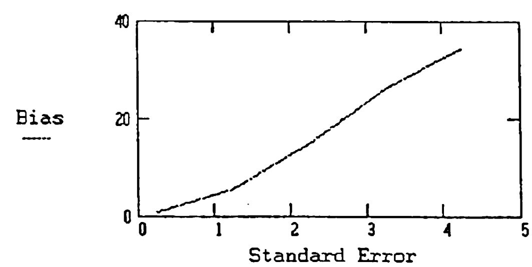

Figure 3: Basis Inflation

The bias increases as the standard error increases, a bias of 20 equates to 10% of z^*, 40 to 20% of z^*. The standard deviation of $\hat{z}^*$ follows a similar trend as shown in the Figure 4.

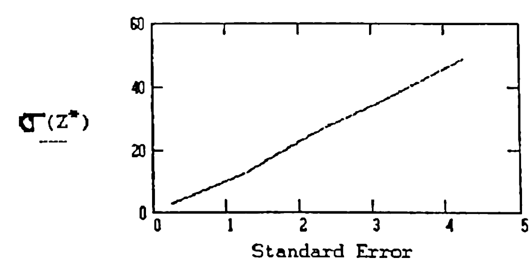

Figure 4: $\sigma(z^*)$ Inflation

Hence, one can expect as σ(parameter estimates) increases not only does the bias increase, but the variance in $\hat{z}^*$ increases.

Figures 5 and 6 take a different view of the problem. In these figures we plot the objective function values on the vertical axis for 8000 samples of the objective function coefficients, and then sort them in descending order within each extreme point. That is,

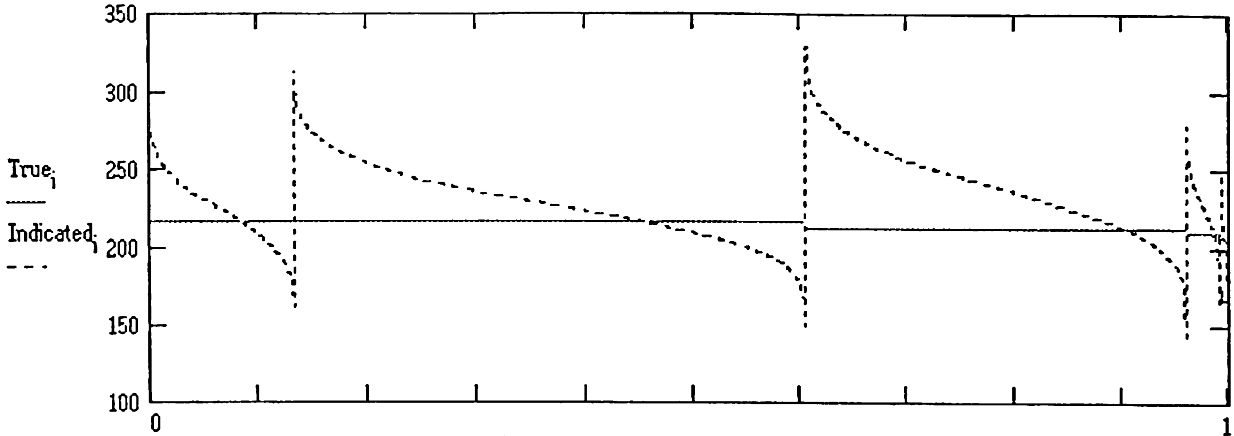

Figure 5: Indicated z* vs. Actual * = 2.25

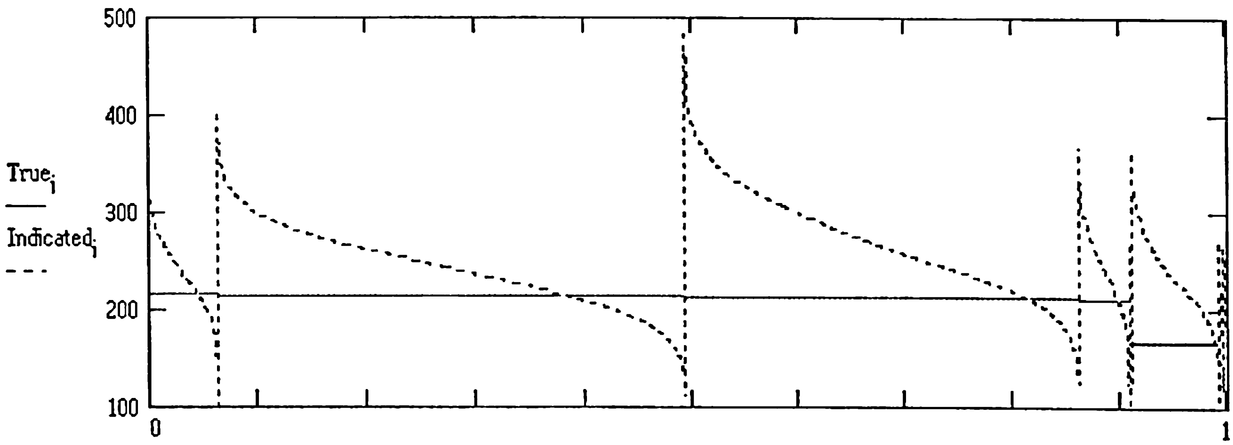

Figure 6: Indicated z* vs. Actual * = 3.25

for a particular sample of **c** we record the objective function value, its chosen extreme point, and the "true" optimal extreme point. The objective function values are then collected and ordered by extreme point.

For any given extreme point the distribution of values are approximately normal. Consider that for a given extreme point (i.e. fixed values of the x_i's),

$$z^* = c_1x_1 + c_2x_2 + c_3x_3 + \ldots + c_nx_n \; . \qquad (6)$$

When examining a particular extreme point the c_i may no longer be normally distributed, since we have a restricted sample of the original c_i values. By Central Limit Theorem arguments, the summation in equation 6 will produce a curve resembling a cumulative normal distribution at each extreme point.

Note that even in choosing an incorrect extreme point (and hence, an inferior strategy), $\hat{z}^*$ can be much higher than even the "true" optimal extreme point z^*. When σ = 2.25 (as in Figure 5), visits to the "true" optimal extreme point occur about 14% of the time, and about 98% of the solutions are "close" to the "true" optimal extreme point. When σ = 3.25 (as in Figure 6) visits to the "true" optimal extreme point occur about 7% of the time, and about 92% of the solutions are "close" to the "true" optimal extreme point.

When σ = 4.25 (figure not shown) visits to the "true" optimal extreme point occur about 6% of the time, and about 91% of the solutions are "close" to the "true" optimal extreme point.

The two principal problems then are the bias in the estimate of the optimal value and the difficulty in identifying the correct or optimal strategy (extreme point). Other techniques must be investigated to put this information into practical application.

4 SAMPLING EXTREME POINTS

We can obtain limited information about the true solution from a single realization of the process. The estimated extreme point may lead to a highly biased solution compared to the true extreme point. Step two shows how to obtain the true extreme point using two methods. The first method samples the generated objective function (in a Monte Carlo fashion) using the variance-covariance matrix from the regression analysis and catalogs the extreme points visited. The second method samples the generated objective function through a design and catalogs the extreme points visited.

As a starting hypothesis, this research assumes that given an initial response surface, its associated variance-covariance matrix, and a large enough sample, we will capture the extreme point corresponding to the optimal solution.

In investigating the Monte Carlo approach we generated samples of the objective function in various specified sizes. We chose these sizes to correspond to the sample sizes required in the experimental design alternative so we can compare the success rate for each level of effort. Each sample objective function inserted into the linear programming problem either yields the true extreme point or not. Obviously the success rate improves as the sample size grows. Repetitively generating these samples enables us to estimate the percentage of times the true solution is captured in a sample of a given size. We generated all samples from the (sampled) multivariate normal distribution describing the response surface (objective function). Table 1 shows the results of the Monte Carlo sampling.

Table 1: Monte Carlo Samples

Standard Error	.25	1.25	2.25	3.25	4.25
N = 49					
% miss "true"	1.2	4.0	5.8	12.3	18.6
N = 74					
% miss "true"	1.0	2.2	4.3	7.7	12.5
N = 100					
% miss "true"	.8	1.8	2.6	5.2	10.0
N = 200					
% miss "true"	1.0	1.8	2.6	5.2	10.0
N = 300					
% miss "true"	.8	1.3	1.5	2.3	3.7
N = 500					
% miss "true"	.5	.9	1.3	1.9	3.1

An advantage to Monte Carlo sampling is that if the analyst has the time and resources, and wants to be conservative, the size of Monte Carlo samples could be increased indefinitely. While Monte Carlo sampling is the least efficient of the options presented here, if the analyst is willing to take enough samples it could be the most effective in sampling the "true" optimal extreme point.

We chose an experimental design as an alternative to the Monte Carlo approach. We investigated the application of a Box-Behnken design from Box and Draper (1987) to the space of **c**. The goal is to capture that true extreme point by systematically investigating the region around our best estimate of the objective function coefficient, $\hat{\mathbf{c}}$.

First we used a modified (only one sample at the zero level) Box-Behnken design. This procedure proved somewhat effective. The sampling was done by varying the estimated objective function coefficients by a percentage of their estimated standard deviation (called standard deviation multiplier) in a method prescribed by the design. This approach was reasonable because we used an orthogonal design to sample the original "Black Box Simulation" to estimate the response surface, as a result, there are no off-diagonal elements in the variance-covariance matrix. A more complicated method may be suggested if off-diagonal elements were present, but it seems an initial orthogonal design is a reasonable approach. Tests using a single Box-Behnken design showed limited success. It appears this single design is inadequate to sample the "true" extreme point.

The second modification to the standard Box-Behnken design was to double the length of the design by sampling at each design point twice. For every identical pair of design points different standard deviation multipliers were used. In effect, a three-level design was transformed into a pseudo five-level design. It is not a true five-level design because each design point has only three levels. It is really the same design run twice with two different standard deviation multipliers. Results with the double Box-Behnken design are superior to sampling in a Monte Carlo fashion 49 times (see Table 2). Results over a broad range of problems indicate this design is superior, but not dramatically, to an equivalent number of Monte Carlo samples. In general, either case fails to give confidence in the results.

The next modification includes adding a third Box-Behnken design to the previous two designs and sampling it at a different standard deviation--this is a pseudo seven-level design. In essence, this is equivalent to sampling from three consecutive designs.

Table 2: Box-Behnken Samples

Standard Error	.25	1.25	2.25	3.25	4.25
Standard Dev (Single)	**2.5**				
% Miss "true"	.2	.8	11.5	32.7	41.6
Standard Dev (Double)	**1.5,**	**3.0**			
% Miss "true"	.25	0.0	1.65	9.1	20.3
Standard Dev (Triple)	**.5,**	**1.75,**	**3.0**		
% Miss "true"	0.1	0.2	0.8	5.0	10.6

Table 3: 5-Level Box-Behnken Type Design

Standard Error	.25	1.25	2.25	3.25	4.25
Standard Dev Single	**2.5**				
% Miss "true"	.5	1.8	2.1	5.3	11.2
Standard Dev Double	**1.5,**	3.0			
% Miss "true"	.2	.3	1.0	2.6	6.8
Standard Dev Triple	**1.5,**	2.75 ,	4.0		
% Miss "true"	0	0	0	.5	2.2

The triple Box-Behnken design had good results, but required more samples. In this case, the triple Box-Behnken design (with four decision variables) required 74 design points. This can be compared to Table 1 with the Monte Carlo experiment of sample 74.

Again there is an advantage to the design over the equivalent number of Monte Carlo samples. The main advantage to a Monte Carlo approach is that the number of samples can be arbitrarily increased to achieve the confidence desired, this may be desirable if a higher confidence in the solution in needed than is possible with this design. To this point, each design was an improvement over an equivalent number of Monte Carlo samples, but no design gave a high success rate at higher noise levels.

In an effort to improve the success rate with higher levels of noise we investigated another type of modification to the basic Box-Behnken design. In this case, we modified the basic structure at each design point. Instead of sampling at the design points using a three-level approach of 1, -1, or 0, this new design was a true five-level design where each design point was sampled with some combination of 1, -1, .5, -.5, or 0. This modification doubles the length of the design and at each design point alternatively samples from either 1 or .5.

The single modified 5-level Box-Behnken design has 49 design points, the same number as the double Box-Behnken design presented in table 2. The 5-level design has a higher success rate in sampling the "true" optimal extreme point than either the double Box-Behnken design, or an equivalent number of Monte Carlo samples. The 5-level Box-Behnken design represents an improvement when sampling at higher noise levels, but the errors could still be considered significant, see Table 3.

A further modification attempts to decrease the errors in sampling the "true" optimal extreme point by doubling the design and choosing a different standard deviation multiplier for the second half of the design. This modification is analogous to the change creating the double Box-Behnken design. This design creates a pseudo nine-level design. Table 3 contains the results of 1000 replications of this design.

The double modified 5-level Box-Behnken design gave excellent results. This design gave the best results for methods with about 97 samples, and it is competitive with a Monte Carlo method of 200 samples.

In the next modification another modified 5-level design is added and sampled at a different standard deviation. This pseudo 13-level design (four variables) has 145 design points. The results in Table 3 show the results are excellent.

The triple 5-level Box-Behnken design was superior to all other designs and even superior to 500 Monte Carlo samples. This design provides excellent sampling in a relatively efficient manner. The main drawback is that it requires 145 samples with only four variables.

Another interesting consideration is the number of extreme points visited with different sampling techniques. If one sampling method provided high accuracy, but required more extreme points to be sampled, then it might not be the best design to employ. Fortunately, no design greatly increased the number of extreme points sampled. Table 4 illustrates the total unique extreme points sampled for 200 Monte Carlo samples and two 5-Level designs -- results are typical of all sampling options.

Table 4: Total Unique Extreme Points

Standard Error	.25	1.25	2.25	3.25	4.25
200 Monte Carlo					
# unique ext. points	3	4	6	8	9
Double 5-level Box-Behnken Type	**1.5,**	**2.5**			
# unique ext. points	2	3	5	8	8
Triple 5-level Box-Behnken Type	**1.5,**	**2.75**	**,4.0**		
# unique ext. points	2	4	5	7	8

5 Screening Extreme Points

In this research, only the objective function is stochastic and therefore only optimality, and not feasibility, is an issue. Using previously recorded solutions we can evaluate new samples to decide whether the linear program needs to be solved. The optimality condition, for a maximization problem, in the general case is:

$$\hat{c} - \hat{c}_B B^{-1} A \leq 0 \qquad (7)$$

where

$\hat{c}$ = estimated objective function
$\hat{c}_B$ = estimated coefficients of the basic variables
B^{-1} = basis inverse
A = constraint matrix

As new extreme points are sampled their corresponding basis inverses are stored and used to screen new

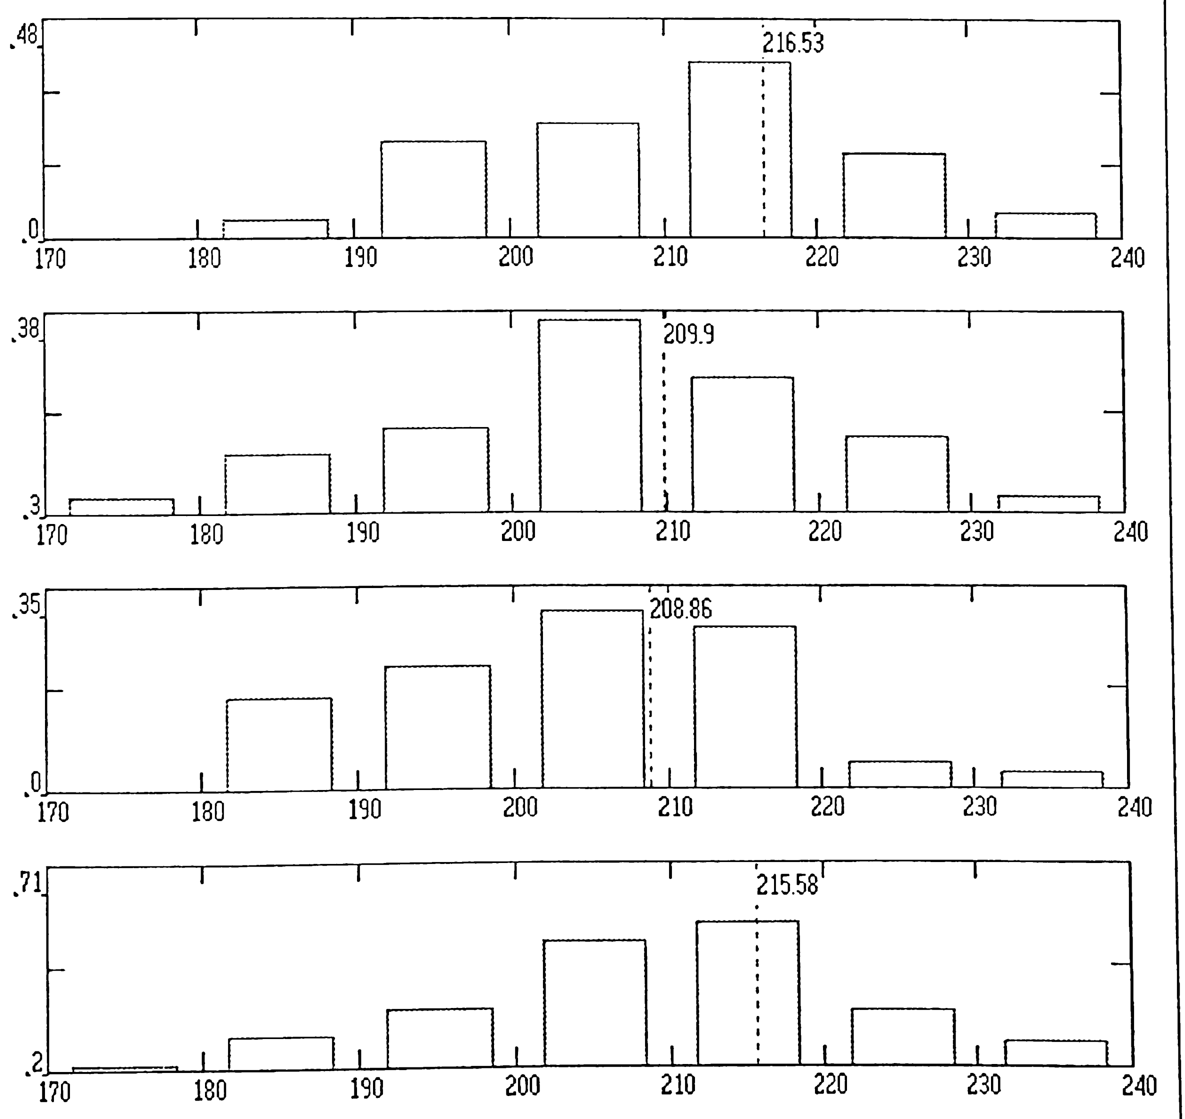

Figure 6: Sample Case Histogram Comparison

objective function samples. For every new sample of the objective function we cycle through Equation 7 until we satisfy the optimality condition. If the optimality condition is never satisfied we then solve the linear program to identify a new basis. Using this scheme we solve a linear program only once for each unique extreme point sampled. The improved efficiency will vary from problem to problem and will also depend on the number of objective function samples, but improvement can be measured in orders of magnitude. Applying this technique greatly increases the practicality and efficiency of the technique. Using this screening procedure makes a strong case for using the triple 5-level Box-Behnken design approach with a large number of samples.

6 SELECTING THE OPTIMAL EXTREME POINT

After identifying the feasible extreme points, we no longer need the linear program, and the decision variable settings at any extreme point are used as input to the simulation to estimate z^*. Once we select a set of decision variable settings, we move to Step 3 in the solution process and the problem becomes selecting the "best" option.

Law and Kelton (1991) present ranking and selection procedures that offer an alternative to the brute force method (independently sampling extreme points) presented above. This ranking and selection procedure was used to analyze the extreme points sampled by the double Box-Behnken type design when the standard error equals 3.25.

After performing a ranking and selection procedure we plot histograms using all simulation samples from the best m alternatives. The histogram can the aid the decision maker by visually representing the possible realizations of the process at given settings. Two important advantages are: avoiding risk by choosing the smallest variance, and illustrating nearly equivalent alternatives and allowing the decision maker to consider factors not captured by the model. A visual representation presents the decision maker with a broader knowledge base from which to make a decision. The histogram is a way to aid the decision maker. In this example, all the actual variances are equal, but the true strength of this method is when the variances are different. Figure 6 illustrates the histograms of the top four alternatives -- dotted vertical lines represent the estimated mean for each alternative. An alternative to presenting a histogram of the data is to plot the normal probability curve defined by the estimated mean and variance. At this point the choice is up to the decision maker.

7 CONCLUSIONS AND RECOMMENDATIONS

The results of this research clearly lead to the conclusion that some kind of variance reduction techniques applied to the simulation would greatly benefit the analyst. If the analyst chooses to use the traditional method of solving this kind of problem (with only one realization of the process) variance reduction procedures appear to be critical if he hopes to have any confidence in the solution. If the analyst chooses to follow the approach recommended in this research variance reduction will play a key role in minimizing the number of extreme points sampled and aiding in the comparison between competing extreme points.

Please refer to Law and Kelton (1991) for explanation of how to apply variance reduction techniques. Some techniques that may be appropriate here are: multiple replications, common random numbers, antithetic random numbers, and control variates.

This paper offers an alternative method to the traditional approach of estimating a response surface and then using it as the objective function of a linear program. On the average the traditional approach will overestimate the true mean response, and it is unlikely we will choose the "true" optimal extreme point. Variance in the estimates of the response surface coefficients can lead to large variance in the estimation of z^* and a low probability of choosing the correct optimal extreme point EP^*. By using the screening procedure this general procedure may become practical for general application.

REFERENCES

Biles, William E. and James J. Swain. "Strategies for Optimization of Multiple-Response Simulation Models," *Proceedings of the Winter Simulation Conference* .135-142. New York: IEEE Press, 1977.

Box, George E. P. and Norman R. Draper. *Empirical Model-Building and Response Surfaces.* New York: John Wiley & Sons, 1987.

Harvey, R. Garrision.. *Optimization of Stochastic Response Surfaces Subject to Constraints with Linear Programming* MS thesis, Air Force Institute of Technology, Wright Patterson Air Force Base, March 1992. (AFIT/GOR/ENS/93M-14)

Law, Averill M. and W. David Kelton. *Simulation Modeling & Analysis.* New York: McGraw-Hill, Inc., 1991.

Morben, Darrel M. *A Simulation Study of an Optimization Model for Surface Nuclear Accession*

Planning. MS thesis, Navel Postgraduate School, Monterey California, September 1989 (AD-A219 228).

Myers, Raymond H., Andre I. Khuri and Walter Carter. "Response Surface Methodology: 1966-1988," *Technometrics, 31*: 137-157 (May 1989).

AUTHOR BIOGRAPHIES

R. GARRISON HARVEY is an analyst with the U.S. Air Force working with the Command Analysis Group at Scott Air Force Base Illinois. He received a B.S. degree in mathematics from the U.S. Air Force Academy in 1988, and he received a M.S. degree in Operations Research form the Air Force Institute of Technology. His research interests are in optimizing airlift systems through simulation and statistical analysis. His mailing address is HQ AMC/XPYP, SCOTT AFB, IL, 62225-5307.

KENNETH W. BAUER JR is an Assistant Professor in the Department of Operational Sciences, School of Engineering, at the Air Force Institute of Technology. He received a B.S. in mathematics from Miami University (Ohio) in 1976, a M.E.A. from the University of Utah in 1980, a M.S. from A.F.I.T. in 1981 and a Ph.D. from Purdue University in 1987. His research interests are in the statistical aspects of simulation. His mailing address is: AFIT/ENS, WPAFB, OH, 45433-0583.

JOSEPH R. LITKO heads the optimization cell within the Command Analysis Group. He also serves as an Adjunct Professor of Operations Research for the Air Force Institute of Technology. His research interests are in combined simulation-optimization modeling of airlift systems and imbedded optimization heuristics for control of simulation models.