CONTROLLED SEQUENTIAL BIFURCATION: A NEW FACTOR-SCREENING METHOD FOR DISCRETE-EVENT SIMULATION

Hong Wan
Bruce Ankenman
Barry L. Nelson
Department of Industrial Engineering and Management Sciences
Northwestern University
Evanston, IL 60208-3119, U.S.A.

ABSTRACT

Sequential bifurcation (SB) is a screening method that is well suited for simulation experiments; the challenge is to prove the "correctness" of the results. This paper proposes Controlled Sequential Bifurcation (CSB), a procedure that incorporates a two-stage hypothesis-testing approach into SB to control error and power. A detailed algorithm is given, performance is proved and an empirical evaluation is presented.

1 INTRODUCTION

Screening experiments are designed to investigate the controllable factors in an experiment with a view toward eliminating the unimportant ones. According to the sparsity of effects principle, in many cases only a few factors are responsible for most of the response variation (Myers and Montgomery 1995). A good screening procedure should correctly and efficiently identify important factors. This is especially important when the system is complicated and many factors are being considered.

In this paper we focus on factor-screening methods for discrete-event simulations. Simulation experiments are significantly different from physical experiments in that they generally involve a large number of factors, and it is easier to implement sequential procedures because of the relatively low cost of switching among settings. Also, it is possible to implement commom random numbers (CRN) to reduce the variance of estimated effect in simulation experiments.

We concentrate on a specific method called Sequential Bifurcation (SB, Bettonvil and Kleijnen 1997). A sequential design is one in which the design points (factor combinations to be studied) are selected as the experiment results become available. Therefore, as the experiment progresses, insight into factor effects is accumulated and used to select the next design point or group of design points.

SB is a group screening procedure in a series of steps. As with other group screening procedures, it is assumed that the sign of each factor effect is known, so that groups contain only factors with effects of the same sign. In each step, a group of factors is tested for importance. The first step begins with all factors of interest in a single group and tests that group's effect. If the group's effect is important, indicating that at least one factor in the group may have an important effect, then the group is split into two subgroups. The effects of these two subgroups are then tested in subsequent steps and each subgroup is either classified as unimportant or split into two subgroups for further testing. As the experiment proceeds, the groups become smaller until eventually all factors that have not been classified as unimportant are tested individually. This method was first proposed for deterministic computer simulations by Bettonvil and Kleijnen (1997). Later the method was extended to cover stochastic simulations (Cheng 1997). The sequential property of the method makes it well suited for simulation experiments. Examples have shown that the method is highly efficient when important factors are sparse and clustered (Cheng 1997, Bettonvil and Kleijnen 1997), but there is no performance guarantee in the stochastic case.

In this paper we propose a modified SB procedure, called Controlled Sequential Bifurcation (CSB), for stochastic simulations. The contribution of CSB is that it controls the Type I Error and power simultaneously. A two-stage testing procedure is introduced to guarantee the power of each step; and at the same time the step-down property of SB implies Type I Error control for each factor.

The paper is organized as follows: In Section 2 we define the underlying response model that we will use. Section 3 describes the procedure and discusses its performance. Section 4 presents an empirical evaluation comparing CSB to another version of SB designed for stochastic simulation.

2 RESPONSE MODEL

In this section we introduce the underlying response model that will guide our new CSB procedure.

2.1 Main-Effects Model

Suppose that there are K factors of interest with effect coefficients $\tilde{\beta} = \{\tilde{\beta}_1, \tilde{\beta}_2, \ldots, \tilde{\beta}_K\}$. The output of interest from a simulation replication is denoted by Y, and Y is represented by the following metamodel:

$$Y = \tilde{\beta}_0 + \tilde{\beta}_1 z_1 + \tilde{\beta}_2 z_2 + \cdots + \tilde{\beta}_K z_K + \varepsilon \tag{1}$$

which is a multiple linear regression model with K regression variables and main effects only. The setting of the factors, $\mathbf{z} = (z_1, z_2, \ldots, z_K)$, is deterministic and under the control of the experimenter. The error term, ε, on the other hand, is a random variable; in this paper we assume it is a $\mathrm{Nor}(0, \sigma^2(\mathbf{z}))$ random variable where $\sigma^2(\mathbf{z})$ is unknown and may depend on $\mathbf{z}$.

We do not assume that the main-effects model holds across the entire range of the factors $\mathbf{z}$. Rather, we assume that it is a good local approximation for modest deviations from a nominal level, typically the center of the design space.

2.2 Determination of Factor Levels

In practice, when we consider whether a change in the response is worth pursuing, the cost to achieve the change is critical. Similarly, when we compare the importance of two different factors we have to make sure that they are based on the same cost or the comparison has little meaning. By scaling the effect coefficients with respect to the cost of changing the factors' levels we can insure that the results have a useful interpretation. We describe one way to do this here.

Let c_i be the cost per unit change of factor z_i, for $i = 1, 2, \ldots, K$. Further, let $c^* = \max_{i \in \mathcal{D}} c_i$, where $\mathcal{D}$ is the set of indices of all of the factors whose levels can only be changed in discrete units (e.g., number of machines at a workstation, or number of cashiers at the checkout). Let Δ_0 be the minimum change in the expected response for which we would be willing to spend c^*, and let Δ_1 be a change in the expected response that we would not want to miss if it could be achieved for only a cost of c^*. If $\mathcal{D} = \emptyset$, then let (c^*, Δ_0) be such that we are willing to spend c^* for a Δ_0 change in the expected response, and define Δ_1 as before.

Let

$$\delta_i = \begin{cases} c^*/c_i, & i \notin \mathcal{D} \\ \lfloor c^*/c_i \rfloor, & i \in \mathcal{D} \end{cases}$$

which is the maximum change in factor i that can be achieved for a cost of c^*; and let $w_i = \delta_i c_i / c^* \le 1$, which is the fraction of a full-cost move, c^*/c_i, that can actually be made for factor i. If factor i can be changed continuously ($i \notin \mathcal{D}$), or $i \in \mathcal{D}$ but c^*/c_i is an integer, then $w_i = 1$. If $i \in \mathcal{D}$ and c^*/c_i is not an integer, then $w_i < 1$.

For instance, suppose that there are $K = 3$ factors. The level of the first can be changed continuously, but the other two are discrete. If $c_1 = 300$, $c_2 = 400$, and $c_3 = 1000$, then $c^* = 1000$, $\delta_1 = 10/3, \delta_2 = 2$, and $\delta_3 = 1$ giving $w_1 = 1, w_2 = 0.8$ and $w_3 = 1$.

Recall that the main-effects model is

$$Y = \tilde{\beta}_0 + \sum_{i=1}^{K} \tilde{\beta}_i z_i + \varepsilon_i.$$

Let the nominal (low) level of z_i be z_i^0 and let the high level be $z_i^0 + \delta_i$, for $i = 1, 2, \ldots, K$. Define the transformed variables $x_i = w_i(z_i - z_i^0)/\delta_i = (c_i/c^*)(z_i - z_i^0)$. Then Y can be expressed as a linear regression on x_i, $i = 1, 2, \ldots, K$, as

$$Y = \beta_0 + \sum_{i=1}^{K} \beta_i x_i + \varepsilon_i \qquad (2)$$

where the low level of x_i is 0, the high level is w_i, and $\beta_i = \delta_i \tilde{\beta}_i / w_i$, for $i = 1, 2, \ldots, K$. Now each β_i, $i > 0$, has a practical interpretation: it represents the change in the expected response when spending c^* to change the level of factor i, and this change can be compared with Δ_0 and Δ_1 without ambiguity. We assume that the sign of each factor effect is known so that we can set the levels of each factor to have $\beta_i > 0$ for all $i > 0$.

2.3 Objective of the Screening Procedure

In screening experiments, the primary objective is to divide the factors into two groups: those that are unimportant, which we take to mean $\beta_i \le \Delta_0$, and those that are important, meaning $\beta_i > \Delta_0$. Since we can never make these determinations with certainty in a stochastic simulation, we instead pursue a screening procedure that controls the probability of incorrectly classifying each factor. More specifically, for those factors with effects $\le \Delta_0$, we require the procedure to control the Type I Error of declaring them important to be $\le \alpha$; and for those factors with effects $\ge \Delta_1$ we require the procedure to provide power for identifying them as important to be $\ge \gamma$. Here α and γ are user-specified parameters and Δ_0 and Δ_1 are defined as in Section 2.2 with $\Delta_1 \ge \Delta_0$. Those factors whose effects fall between Δ_0 and Δ_1 are also considered important and we want the procedure to have reasonable, though not guaranteed, power to identify them. Figure 1 is a generic illustration of the desired performance of our screening procedure.

To illustrate, consider a simulated manufacturing system where the response is the expected throughput of the system. The controllable factors may include the number of machines at each workstation; average processing time of each machine; and skill levels of the workers. The practical threshold Δ_0 is set as the minimum change in expected throughput that managers consider worth pursuing at a cost c^* of changing the most expensive factor by one unit. For example, c^* might be the cost of purchasing a very expensive machine. In this illustration, screening experiments would be used to identify each factor that influences the expected throughput by more than Δ_0 when spending c^* to change that factor. For each factor, the procedure should have probability $\le \alpha$ of declaring it important if it cannot influence the expected throughput by at least Δ_0 at a cost of c^*. The procedure should also have probability $\ge \gamma$ of identifying a factor as important if its influence on the expected throughput is $\ge \Delta_1$ at a cost of c^*. Here Δ_1 is a critical change in the expected throughput that the managers do not want to ignore if it can be achieved for a cost of only c^*. Factors whose effects are neither unimportant nor critical will be identified with less power than γ.

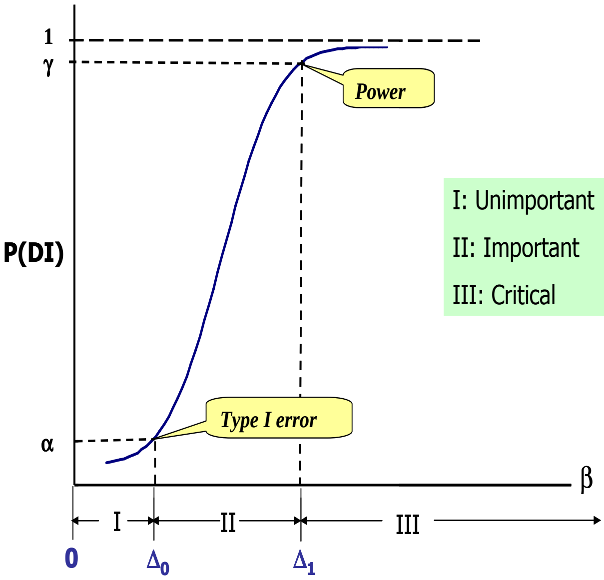

Figure 1: Generic Illustration of Desired Performance of Screening Procedures

3 CONTROLLED SEQUENTIAL BIFURCATION (CSB)

The CSB procedure inherits its basic concepts from the SB procedure proposed by Bettonvil and Kleijnen (1997) and from the SB-under-uncertainty procedure proposed by Cheng (1997). Specifically, like other SB procedures, the CSB procedure is a series of steps in which groups of factors are tested. If a group of factors is considered unimportant, then every factor in the group will be considered unimportant. If the group is considered important, then it is split for further testing. When the algorithm stops, each of the factors will be classified as either important or unimportant. The unique feature of CSB is that each step contains a two-stage testing procedure to insure the desired power. In addition, CSB preserves the step-down nature of SB so that Type I Error is controlled. The testing procedure is explained in detail in the following sections.

3.1 Notation

The notation that we use to define CSB is provided below.

There are in total K indexed factors.

Let x_i represent the setting of factor i. A replication at level k is defined as follows:

$$x_i(k) = \begin{cases} w_i, & i = 1, 2, \ldots, k \\ 0, & i = k+1, k+2, \ldots, K \end{cases}$$

$Y_j(k)$: The j^{th} response at level k

$\bar{Y}(k)$: Average of all available responses at level k

n_0: Number of initial replications made at each level

σ_k^2: variance of responses at level k

$D_j(k_1, k_2) = Y_j(k_2) - Y_j(k_1), j = 1, 2, \ldots,$ for $k_2 > k_1$, whose expected value is $\sum_{i=k_1+1}^{k_2} w_i \beta_i$; and whose variance is $\sigma_{k_1}^2 + \sigma_{k_2}^2$.

$\bar{D}(k_1, k_2) = \bar{Y}(k_2) - \bar{Y}(k_1)$, for $k_2 > k_1$.

$w(k_1, k_2) = \min\{w_{k_1+1}, w_{k_1+2}, \ldots, w_{k_2}\}$ is the smallest weight associated with $\beta_{k_1+1}, \beta_{k_1+2}, \ldots, \beta_{k_2}$.

$S^2(k_1, k_2) = \sum_{j=1}^{n_0} \left(D_j(k_1, k_2) - \bar{D}(k_1, k_2)\right)^2 / (n_0 - 1)$. Notice that $S^2(k_1, k_2)$ is only determined by the initial n_0 replications.

$U_A(k_1, k_2) = \Delta_0 + t_{\sqrt{1-\alpha}, n_0 - 1} S(k_1, k_2)/w(k_1, k_2)\sqrt{n_k}$, where $n_k = \min\{n_{k_1}, n_{k_2}\}$ and n_{k_i} is the total number of available responses at factor level k_i. The subscript $A = I, II$ denotes the first or second stage of the testing procedure, respectively.

$L_A(k_1, k_2) = \Delta_0 - t_{(1+\gamma)/2, n_0 - 1} S(k_1, k_2)/w(k_1, k_2)\sqrt{n_k}$, where $n_k = \min\{n_{k_1}, n_{k_2}\}$ and n_{k_i} is the total number of available responses at factor level k_i. The subscript $A = I, II$ denotes the first or second stage of the testing procedure, respectively.

h: A constant such that $\Pr(T \leq t_{\sqrt{1-\alpha}, n_0} - h) = (1-\gamma)/2$, where T is a t-distributed random variable with $n_0 - 1$ degrees of freedom.

$N(k_1, k_2) = \lceil h^2 S^2(k_1, k_2)/w^2(k_1, k_2)(\Delta_1 - \Delta_0)^2 \rceil$

3.2 CSB Procedure

A high-level description of CSB is shown in Figure 2. The figure illustrates how groups are created, manipulated, tested and classified, but does not specify how data are generated or what tests are performed. Detailed descriptions of data collection and hypothesis testing follow. This section is closed by an example.

Data (replications) are obtained whenever new groups are formed according to the following rule: When forming a new group containing factors $\{k_1 + 1, k_1 + 2, \ldots, k_2\}$ with $k_1 < k_2$, check the number of observations at level k_1 and k_2.

If $n_{k_1} = 0$, then get n_0 observations at level k_1 and set $n_{k_1} = n_0$.

If $n_{k_2} = 0$, then get n_0 observations at level k_2 and set $n_{k_2} = n_0$.

If $n_{k_1} < n_{k_2}$, then make $n_{k_2} - n_{k_1}$ additional replications at level k_1 and set $n_{k_1} = n_{k_2}$.

If $n_{k_2} < n_{k_1}$, then make $n_{k_1} - n_{k_2}$ additional replications at level k_2 and set $n_{k_2} = n_{k_1}$.

Suppose the group removed from the queue contains factors $\{k_1 + 1, k_1 + 2, \ldots, k_2\}$ with $k_1 < k_2$. The **Test** step in Figure 2 tests the following hypothesis to determine if a group might contain important factors:

$$H_0 : \sum_{i=k_1+1}^{k_2} \beta_i \leq \Delta_0 \ vs. \ H_1 : \sum_{i=k_1+1}^{k_2} \beta_i > \Delta_0.$$

Initialization: Create an empty LIFO queue for groups. Add the group $\{1, 2, \ldots, K\}$ to the LIFO queue.

While queue is not empty, do

Remove: Remove a group from the queue.

Test:

Unimportant: If group is unimportant, then classify all factors in the group as unimportant.

Important (size = 1): If group is important and of size 1, then classify the factor as important.

Important (size > 1): If group is important and size is greater than 1, then split it into two subgroups such that all factors in the first subgroup have smaller index than those in the second subgroup. Add each subgroup to the LIFO queue.

End Test

End While

Figure 2: Structure of CSB

The procedure given below for testing this hypothesis guarantees power $\geq \gamma$ if $\sum_{i=k_1+1}^{k_2} \beta_i \geq \Delta_1$.

1. If $\bar{D}(k_1, k_2)/w(k_1, k_2) \leq U_I$, and $\min\{n_{k_1}, n_{k_2}\} \geq N(k_1, k_2)$, then classify the group as unimportant.

2. Else if $\bar{D}(k_1, k_2)/w(k_1, k_2) \leq L_I$, then classify the group as unimportant.

3. Else if $\bar{D}(k_1, k_2)/w(k_1, k_2) > U_I$, then classify the group as important.

4. Else make $(N(k_1, k_2) - n_{k_1})^+$ observations at levels k_1 and k_2 (recall that $n_{k_1} = n_{k_2}$). Then set $n_{k_1} = n_{k_2} = \max\{N(k_1, k_2), n_{k_1}\}$. Notice that $S^2(k_2, k_2)$ and the degrees of freedom do not change, but $\bar{D}(k_1, k_2)$ is updated.

   (a) If $\bar{D}(k_1, k_2)/w(k_1, k_2) < U_{II}$, then classify the group as unimportant.

   (b) If $\bar{D}(k_1, k_2)/w(k_1, k_2) \geq U_{II}$, then classify the group as important.

Notice that $E\left[\bar{D}(k_1, k_2)\right] = \sum_{i=k_1+1}^{k_2} w_i \beta_i \leq \sum_{i=k_1+1}^{k_2} \beta_i$. Therefore testing based on $\bar{D}(k_1, k_2)$ would sacrifice power. Thus, we use $\bar{D}(k_1, k_2)/w(k_1, k_2)$ because $E\left[\bar{D}(k_1, k_2)/w(k_1, k_2)\right] \geq \sum_{i=k_1+1}^{k_2} \beta_i$.

As an illustration, consider the case of $K = 10$ factors and the first pass through the algorithm. Initially we make n_0 replications at level 0 (all factors at their low level) and n_0 replications at level 10 (all factors at their high level). The group removed from the queue contains all factors and $w(0, 10) = \min\{w_1, w_2, \ldots, w_{10}\}$.

Next we evaluate $\bar{D}(0, 10)$, U_I and L_I. If $\bar{D}(0, 10)/w(0, 10) \leq L_I$, then we conclude that none of the factors are important, since the sum of all effects is not important, and the algorithm stops. If $\bar{D}(0, 10)/w(0, 10) > U_I$, then the factors are separated into two groups, $\{\beta_1, \beta_2, \beta_3, \beta_4, \beta_5\}$ and $\{\beta_6, \beta_7, \beta_8, \beta_9, \beta_{10}\}$, and n_0 replications are made at level 5 (x_i, $i = 1, 2, \ldots, 5$ are set at their high level and x_i, $i = 6, 7, \ldots, 10$ are set at their low level). Both groups are added to the queue.

If, on the other hand, $\bar{D}(0, 10)/w(0, 10)$ is between L_I and U_I, then we calculate $N(0, 10)$. If $N(0, 10) \leq n_0$, then we conclude that all the factors are not important and the algorithm stops. If $N(0, 10) > n_0$, then we collect $N(0, 10) - n_0$ replications at both level 0 and level 10, reevaluate $\bar{D}(0, 10)$, and calculate U_{II}. If $\bar{D}(0, 10)/w(0, 10) \geq U_{II}$, then the factors are separated into two groups as described above and n_0 replications are made at level 5. Both groups are added to the queue. Otherwise, all factors will be considered as unimportant and the algorithm stops.

3.3 Performance of CSB

The performance guarantees for the CSB procedure are stated in following theorems. For the proofs see Wan, Ankenman and Nelson (2003).

Theorem 1 *If model (2) holds with normally distributed error, then CSB guarantees that*

$$\Pr\{\textit{Declare factor } i \textit{ important } |\beta_i \leq \Delta_0\} \leq \alpha$$

for each factor i individually.

Theorem 2 *Let the group containing the factors denoted* $\{k_l + 1, \ldots, k_m\}$ *be represented by* $\{k_l \to k_m\}$, $0 \leq k_l \leq k_m \leq K$. *If model (2) holds with normally distributed error, then the two-stage test guarantees that*

$$\Pr\left\{\textit{Declare } \{k_l \to k_m\} \textit{ important } \left| \sum_{i=k_l+1}^{k_m} \beta_i \geq \Delta_1 \right.\right\} \geq \gamma$$

for each group $\{k_l \to k_m\}$ *tested.*

In summary, the CSB procedure controls the Type I Error for each factor individually and guarantees the power for each step. The procedure does not require an equal-variance assumption, and is valid with or without common random numbers. The empirical evaluation will be discussed in Section 4.

4 EMPIRICAL EVALUATION

In this section, we discuss the numerical results of simulation experiments to compare the following two procedures:

1. The CSB method proposed in Section 3.

2. Cheng's method (Cheng 1997), an enhancement of the SB procedure for stochastic responses that assumes equal variances.

The idea behind Cheng's method is to determine whether a group of two or more factors are unimportant by constructing a one-sided confidence interval on the group effect. For a group containing a single factor, replications are added one-at-a-time until a two-sided confidence interval on the factor effect shows that the effect is important or unimportant. When a single factor is tested, the method employs an indifference parameter a. In our notation, all the factors with effects smaller than $\Delta_0 + a$ can be classified as unimportant. Cheng's method does not guarantee to control Type I Error for each factor or power at any step, and has no concept like Δ_1 for a critically important factor.

4.1 Summary of Results

Rather than employ system simulation models in this test, we chose instead to generate data from a main-effects model in which we could control the size of the effects and the variances at different design points. Normal errors are assumed with mean 0 and standard deviation, σ, equal to $m * (1 + \mathcal{I} * \text{size of the group effect})$, where $\mathcal{I}$ is 0 if we are running an equal-variance case, and 1 for an unequal-variance case. Thus, in unequal variance cases the standard deviation is proportional to the size of the effect of the group being screened. Neither procedure assumes prior knowledge of the variances. Common random numbers were not employed.

For each case considered, the CSB procedure is applied 1000 times and the percentage of time factor i is declared important is recorded; this is an unbiased estimator of Pr{factor i is declared important}.

To compare CSB to Cheng's method, we set the indifference parameter, a, such that the number of replications required by Cheng's method is approximately the same as the number used by CSB for that case. Therefore we can compare the achieved Type I error and power of the two methods with equal simulation effort.

The performance of Cheng's method depends on the case considered. When the variances are large or unequal, Cheng's method loses control of the Type I Error and power. The CSB method, on the other hand, controls the Type I Error and power across all cases (although the number of replications required to achieve this does differ substantially by case).

In the following subsection we provide some illustrative numerical results that emphasize the key conclusions.

4.2 Unequal-Variance Cases

We set the parameters as in Table 1. We considered two different settings for the factor effects:

1. In Case 1 we set $(\beta_1, \beta_2, \ldots, \beta_{10}) = (2, 2.44, 2.88, 3.32, 3.76, 4.2, 4.64, 5.08, 5.52, 6)$, spanning the range from Δ_0 to $\Delta_0 + \Delta_1$. For CSB, the probability that β_1 is declared important should be smaller than 0.05, but for $\beta_6, \ldots, \beta_{10}$ it should be ≥ 0.95.

   Letting $P(DI)$ mean "probability of being declared important," Figure 3 plots $P(DI)$ against effect size for Cheng's method and CSB with large ($m = 1$) and small ($m = 0.1$) variances. We can see that when variance is small, the two methods have similar performance although CSB attains greater power earlier. When the variance is large, however, Cheng's method loses control of both Type I Error and power.

Table 1: Parameters for Unequal-Variance Cases

Parameter	Value
K	10
Δ_0	2
Δ_1	4
α	0.05
γ	0.95
σ	$m*(1 + \text{size of the group effect})$
m	0.1, 1

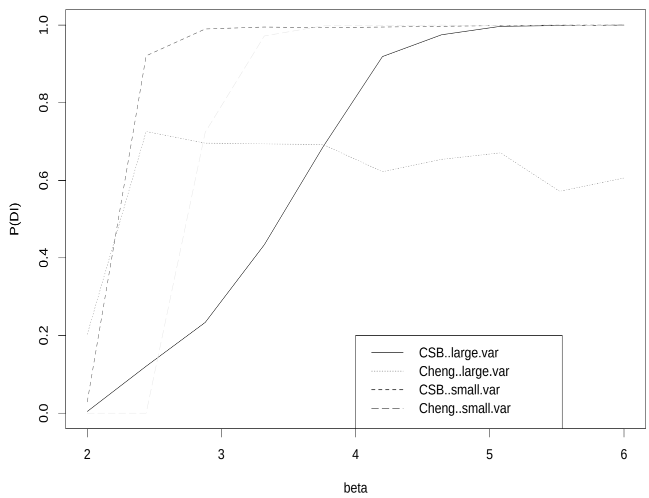

Figure 3: Case 1 with Unequal Variances

2. In Case 2 we set $(\beta_1, \beta_2, \ldots, \beta_{10}) = (2, 2, 2, 2, 2, 2, 2, 2, 2, 2)$, so that all effects are Δ_0. This set is designed to study the Type I Error control of the two methods. The other parameters are the same as in the previous case.

   Figure 4 shows the Type I Error control of both methods. Cheng's method has large Type I Error (as high as 0.5) when the variance is large. Even for the small-variance case, the largest Type I Error is still more than 0.2 for Cheng's method. By design, CSB controls Type I Error to be $\leq \alpha$ in all cases.

To summarize, CSB has superior performance relative to Cheng's method in large and unequal variance cases. CSB has guaranteed performance with different parameter and factor configurations, which makes it attractive for problems with limited prior knowledge. Cheng's method, on the other hand, assumes variance homogeneity to gain advantages on degrees of freedom and it can be effective when this assumption is satisfied.

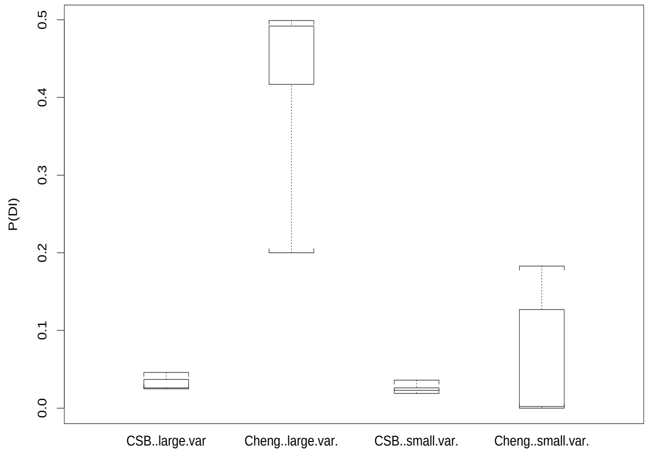

Figure 4: Case 2 with Unequal Variances

5 CONCLUSION

CSB is a new factor-screening method for discrete-event simulations; it combines a two-stage hypothesis-testing procedure with the sequential bifurcation method to control the power at each bifurcation step and Type I Error for each factor under heterogeneous variance conditions. CSB is the first factor-screening procedure to provide these guarantees.

ACKNOWLEDGMENT

This project was partially supported by a grant from General Motors R&D.
Partions of this work is published in *Proceedings of the 2003 Winter Simulation Conference*, (S. Chick, P. J. Sánchez, D. Ferrin, and D. J. Morrice, eds.) Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

REFERENCES

Bettonvil, B., and J. P. C. Kleijnen. 1997. Searching for important factors in simulation models with many factors: Sequential bifurcation. *European Journal of Operational Research* 96 (1): 180–194.

Cheng, R. C. H. 1997. Searching for important factors: Sequential bifurcation under uncertainty. In *Proceeding of the 1997 Winter Simulation Conference*, ed. S. Andradótir, K. J. Healy, D. H. Withers, and B. L. Nelson, 275–280. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers. Available online via `<http://www.informs-cs.org/wsc97papers/0275.PDF>` [accessed July 6, 2003].

Myers, R. H., and D. C. Montgomery. 1995. *Response Surface Methodology: Process and Product Optimization Using Designed Experiments*. New York: John Wiley & Sons Inc.

Wan, H., B. Ankenman, and B. L. Nelson. 2003. Controlled sequential bifurcation: A new factor-screening method for discrete-event simulation. Technical Report. Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL.