QUALITY AND RELIABILITY ENGINEERING INTERNATIONAL
Qual. Reliab. Engng. Int. 2000; **16**: 355–362

COMPARING EXPERIMENTAL DESIGN STRATEGIES FOR QUALITY IMPROVEMENT WITH MINIMAL CHANGES TO FACTOR LEVELS

WILLIAM R. MCDANIEL AND BRUCE E. ANKENMAN*
Department of Industrial Engineering and Management Sciences, Northwestern University, 2145 Sheridan Road, Tech C210, Evanston, IL 60208-3119, USA

SUMMARY
The 'small factor change' problem, where an experimental design strategy is used to find a certain amount of improvement in a response while changing the factor levels as little as possible, is addressed. Using a recently developed test bed for response surfaces, we have simulated a broad range of response surface functions and collected empirical results on the performance of seven experimental design strategies when confronted with this problem. Copyright © 2000 John Wiley & Sons, Ltd.

KEY WORDS: small factor change; experimental design; response surface methodology

1. INTRODUCTION

In this research, a set of experimental design strategies is applied to a situation that we call the *small factor change* problem to determine which of these strategies performs best on selected measures. The goal of experimentation in the small factor change problem is to gain a specific amount of improvement in a response while changing the factor levels as little as possible. As an example, consider an automobile design problem where there is a specified miles per gallon (*MPG*) rating desired. Some of the primary factors that may affect the *MPG* are the aerodynamics of the car and the composition of the materials used in manufacture. However, making changes to either of these factors may have an effect on the appearance of the automobile. For some automobiles, such as the Ford Mustang or Chevrolet Corvette, the appearance is very distinctive and is an important selling point, so very little change can be made to factors affecting the appearance. The experimental problem is to find a specific improvement in the response, the *MPG*, while changing the factors, the aerodynamics and materials, as little as possible.

The small factor change problem can be presented more formally as an experimental design problem by first assuming that there are k factors of interest, $\mathbf{x} = (x_1, x_2, \ldots, x_k)'$ and an observed response y where $y = f(\mathbf{x}) + \varepsilon$, and ε is a random error term such that $E(\varepsilon) = 0$. Let the current settings of the factors be represented by $\mathbf{x}^* = (x_1^*, x_2^*, \ldots, x_k^*)'$. In the small factor change problem, changing any factor is presumed to have a detrimental effect on some characteristic other than y. For convenience, assume that the factors are scaled such that a change of one unit in any of the factors can be expected *a priori* to have the same detrimental effect as a change of one unit in any other factor. If this is not the case, then weights can be introduced to scale the factors to best represent the severity of changing each factor. The objective of the small factor change problem is to find the solution point $x^{\text{opt}} = (x_1^{\text{opt}}, x_2^{\text{opt}}, \ldots, x_k^{\text{opt}})$ satisfying

$$\text{Min} \sum_{i=1}^{k} |x_i^* - x_i^{\text{opt}}|$$
$$\text{Subject to: } E(y \mid \mathbf{x}^{\text{opt}}) = y^{\text{opt}}$$

where y^{opt} is the desired mean response. While the constraint is given as an equality, it may also be represented by an inequality, depending on the nature of the target (smaller the better, larger the better, or nominal the best).

Despite the optimization framework, the small factor change problem is an experimental design problem because the function, f, is not generally known and must be approximated by fitting some

*Correspondence to: B. E. Ankenman, Department of Industrial Engineering, Northwestern University, 2145 Sheridan Rd, Evanston, IL 60208-3119, USA.

model to collected data. Generally, the number of experimental observations is constrained by resource issues. Thus, the experimental design problem is to choose design points (factor settings) for collecting data that will allow for approximating the function f with a fitted model g as accurately as possible in a region that includes y^{opt}, so that x^{opt} can be accurately estimated.

When searching for an optimum, it is advantageous to use fully sequential experimental designs, where each design point is selected based on the information obtained from the prior designs points. To achieve some of the advantages of a sequential design, but to reduce complexity, experimental designs are typically run in phases where a set of design points is selected and then, based on the analysis of those data, another set of design points is selected. We will call these partially sequential designs, *experimental design strategies*.

There are several experimental design strategies that may be considered for the small factor change problem. One example is traditional response surface methodology which combines an initial screening design with a line search in the steepest direction of ascent followed by a capping design, such as a central composite (see Myers and Montgomery, [1]). Because changing many factors is not desirable in the small factor change problem, some variation on a one-factor-at-a-time strategy may be appropriate. Generally, statisticians have shunned this strategy because it does not take into account the possibility of interactions between factors and is often less efficient than other methods. However, strategies involving at least some one-factor-at-a-time exploration may hold promise in the small factor change problem because they can detect significant main effects and may allow for fewer factor changes.

In an experimental design strategy, decisions concerning the next phase of the design are dependent on the response surface function. The effectiveness of the strategies cannot be determined easily unless the specific response surface function is known. Since the response surface function is rarely known, design strategies do not lend themselves well to analytic study and comparisons between strategies are usually not available. In this article, we use a simulation study to collect empirical data on the ability of different experimental design strategies to solve the small factor change problem. In McDaniel and Ankenman [2], a response surface test bed was proposed that would produce random response surface functions while allowing for some control of the characteristics of the surfaces. Using this test bed, we have produced surfaces with a broad range of different characteristics to determine which design strategies perform well on the small factor change problem over a wide variety of response surface types.

In Section 2, we present the motivating example for this research. In Section 3, our test bed of simulated response surfaces is briefly described. In Sections 4 and 5, the simulation study used to investigate this problem is described. Section 6 presents the results of the simulation study, where a form of traditional response surface methodology was shown to perform better than any other design strategy that we tested.

2. MOTIVATING EXAMPLE

In the example that motivated this work, transmission fluid was formulated and subjected to a battery of expensive tests for attributes such as wear resistance, thermal breakdown and viscosity. A formulation was found that met all of the requirements except for the wear resistance specification. Due to the expense of the testing, the design engineers decided to run experiments on a bench-top wear test to alter the formula to improve on the wear reduction response. Since the other responses were not being measured on the altered formulae, the goal was to improve the wear resistance, but minimize the changes in the formula so the results of other tests would not change significantly. When a desirable outcome was achieved for the wear resistance, the whole battery of tests was rerun on the resulting product to confirm that the small changes made in the factor levels did not adversely affect the other responses of interest. Although a fractional factorial design was used to run the bench-top wear experiment, there was some controversy over whether a one-factor-at-a-time design would be more appropriate for finding an improvement with minimal changes. There was also discussion of what follow up designs might be appropriate.

This research is intended to provide guidance for experimental design in small factor change problems which could be encountered in many physical systems, such as printing ink formulations, electronics development, or the manufacture of prepared foods [3].

3. RESPONSE SURFACE TEST BED DESIGN

To study experimental design strategies for the small factor change problem, we created a *test bed* that will simulate different response surfaces using polynomial functions. This test bed, which is described in detail in [2] generates these functions randomly while still maintaining some control over their response range, their dimensionality and their *bumpiness* (the presence and size of local minimum, maximum, or inflection points). The test bed was designed to control the general characteristics of the surfaces produced, but still maintain enough randomness so that the surfaces in the test bed resemble what might be encountered in physical systems. A brief description of the types of surface characteristics controlled by the test bed user follows.

Included in the test bed is the ability to control effect sparcity and effect heredity. Effect sparcity refers to the conjecture that only a few of several possible factors will have a significant effect on the response. Effect heredity refers to the common assumption that if a two-factor interaction is significant, at least one of the main effects that make it up is also significant [4,5]. The test bed is also designed to give the user some control over the number of higher-order main effect terms, such as x_1^2, and interactions, such as x_1x_2, which in turn control the bumpiness of the surface.

The test bed uses a *region of operability* for each of the factors, which is the range over which experimentation may take place. Because any physical system will have constraints on the possible values of the factor settings, this concept was included in the test bed surfaces. Similarly, within the region of operability of the factors the test bed approximately controls the *response range*, which is the expected range of values of the response on the surface. This provides a sense of scale to the response that would be assumed in a physical system but must be defined for a simulated response surface function. There is also a means of controlling the relative *flatness* of the surfaces created by the test bed, so that local deviations are relatively large or relatively small with respect to the response range. Additionally, the test bed allows the user to control the distribution of the error term applied to the observed responses.

The test bed was implemented and the simulation study described in the next section was run in the S-Plus® programming language on a PC running Windows 95. S-Plus® was chosen because the language has many useful statistical tools and features for both surface creation and data analysis [6].

4. SIMULATION METHODOLOGY

In the simulation study, seven design strategies, which are described in the next section, are compared for a small factor change problem. In order to use a wide variety of surfaces from the test bed, 49 different surface types were created as part of a orthogonal array of input settings for the test bed. The orthogonal array varied seven surface characteristics each with seven levels. A description of the surface characteristics and their levels are shown in Appendix A and copy of the orthogonal array is provided in Appendix B.

For each of the 49 surface types, 5 surface forms were created. A surface form is a symbolic representation of the polynomial response surface. For example, a surface form is a symbolic equation as follows:

$$y = \gamma_0 + \gamma_1^{(1)}x_1 + \gamma_2^{(1)}x_2 + \gamma_3^{(1)}x_3 + \gamma_{12}x_1x_2 + \cdots + \gamma_1^{(2)}x_1^2 + \cdots + \gamma_3^{(6)}x_3^6$$

where $\gamma_j^{(i)}$ refers to coefficient of the ith order main effect term for factor j and γ_{ijk} is the coefficient for the interaction between factors i, j, and k. The non-zero coefficients in the surface form identify the terms that are in the response surface function. A response surface function is created by generating numerical values for the non-zero coefficients in the surface form. For each surface form, five response surface functions were generated. These five functions will have the same active factors and interaction effects, but will have different values for each of the coefficients. In total, there were 1225 response surface functions created (49 surface types × five surface forms × five response surface functions). On each of these response surface functions, five random starting points in the region of operability were selected and all seven design strategies were conducted from each starting point. Thus, a total of 42 875 simulation runs were made for this study (1225 response surface functions × five starting points × seven design strategies).

For each simulation run, the goal was to find a response, v, that is a 10% increase in the response over the value of the response, u, at the starting point (i.e. $v = 1.1u$). The 10% value was chosen arbitrarily, but should have little effect on the qualitative results since the scaling of the problem is relative. There were 15 possible factors that might be active on each response surface function.

Once a design strategy had completed all of it's allotted experimental runs from a given starting point, the data were used to model the local region of the response surface with a second-order polynomial function that we call the *fitted model*. Using this fitted model, a solution for the small factor change problem was estimated by solving a nonlinear optimization problem. Using the nonlinear optimizer routine *nlminb* provided with S-Plus®, the penalty function method (ss [7, p. 361]) was used with an objective function of the form:

$$\text{Min obj} = |x_1 - x_1^*| + |x_2 - x_2^*| + \cdots + |x_{15} - x_{15}^*| + 1000^*(v - \psi(\mathbf{x}, \mathbf{b}))^2$$

where $\mathbf{x}^* = (x_1^*, x_2^*, \ldots, x_{15}^*)$ is the starting point, $\mathbf{x} = (x_1, x_2, \ldots, x_{15})$ is the estimated solution found by the algorithm, and $\psi(\mathbf{x}, \mathbf{b})$ is the estimated response at $\mathbf{x}$ resulting from the fitted model with parameters $\mathbf{b}$ that are estimated by least squares from data collected by the design strategy. This objective function solves for the minimum distance to a point on the fitted model ψ with a response that is 10% larger than the response at the starting point. The penalty factor of 1000 forces the algorithm to search for a solution on the fitted model with less than 1% error between v and $\psi(\mathbf{x}, \mathbf{b})$. Because algorithms for solving nonlinear problems are sensitive to the starting solution provided, it is possible that a better solution exists than that found by this algorithm, however, we treat these solutions as optimal for each ψ.

5. DESIGN STRATEGIES

From each starting point, each of the seven design strategies was allotted 33 experimental trials to estimate the response surface function locally. All design strategies were conducted in two sequential phases, an initial screening design (phase I) followed by a re-scaled experiment in the three most important factors (phase II). The data collected in both phases was to create the fitted model and find an estimated solution point, $\mathbf{x}$ as described above.

In phase I designs, it is assumed that the region of operability for each factor may be much larger than the region of interest required for solving the small factor change problem. The scale for all strategies in the first phase was the same (starting point ±3% of each factor range). The results from the phase I designs were used to re-scale the experimental region for the phase II designs. Only the three most significant factors were used in phase II of each strategy. The experimental region in these three factors was scaled to include a 10% improvement assuming a simple linear model from the phase I data.

A constraint was placed on the scale of the phase II experiment so the experimental region for any factor would not extend to more than 30% of the range of that factor from the starting point. Since small factor change problems often have limits on the amount that a factor can be changed, the solution was also constrained to be within 30% of the factor range from the starting point.

Table 1. Raw data from a single starting point on a surface of type 10

Strategy	True response	Distance
FF–BB	48.11	29.28
FF–CC	48.93	29.97
FF–OFAAT′	46.67	21.15
FF–SC	48.72	35.48
OFAAT–BB	44.91	30.08
OFAAT–CC	44.75	28.15
OFAAT–OFAAT′	46.18	42.29

Table 2. MINITAB® statistics on the average percent difference in the distance to successful solutions for each strategy

Strategy	N	Mean	Standard deviation	SE mean
FF–BB	30	6.188	4.491	0.820
FF–CC	34	4.695	4.554	0.781
FF–OFAAT′	25	13.47	10.58	2.12
FF–SC	32	11.83	9.43	1.67
OFAAT–BB	27	16.10	12.23	2.35
OFAAT–CC	30	8.11	7.41	1.35
OFAAT–OFAAT′	24	25.50	43.07	8.79

The phase I designs were chosen as screening designs to sort through the 15 factors to find the three most significant. Two phase I designs were used: a fractional factorial (FF) design and a one-factor-at-a-time (OFAAT) design with 2 levels per factor. Both of these designs used 16 experimental trials. The phase II designs were chosen to map the factor space near the starting point in the three most significant factors. The four phase II designs are: a Box–Behnken (BB) design, a central composite (CC) design, a one-factor-at-a-time design with five levels per factor plus two center points (OFAAT′), and a sliding cube (SC) design. The SC follow-up design was proposed by Mee [8] and for this study simply meant a full factorial in the three factors shifted in the direction of improvement for each factor. This was combined with additional center points at both the starting point and the center of the SC. Each of these Phase II designs consisted of 17 experimental trials.

A design strategy is defined by a phase I design followed by a phase II design. The phase I–phase II design combinations that made up the seven design strategies are FF–BB, FF–CC, FF–OFAAT′, FF–SC, OFAAT–BB, OFAAT–CC and OFAAT–OFAAT′. Notice that one combination of phase I–phase II designs, the OFAAT–SC strategy, was not used since the first design was not a cube design.

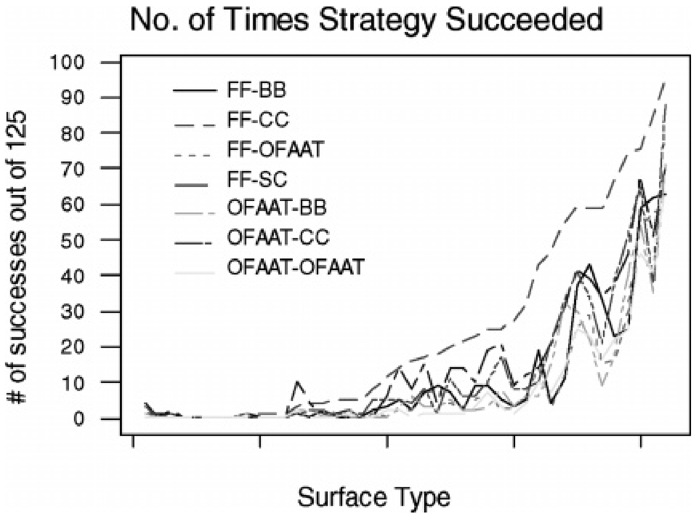

Figure 1. Comparison of the number of successes for each strategy

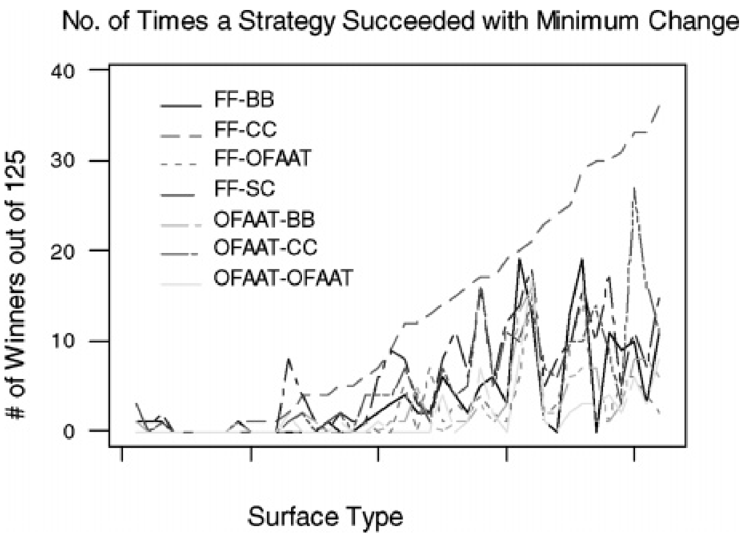

Figure 2. Comparison of the number of times each strategy succeeded with the least change in the factors

6. CRITERIA AND RESULTS

Three metrics were collected for each strategy on each of the 49 surface types so that direct comparisons could be made between the strategies on a given type of surface. The first metric, which we call *success*, is a tally of the number of times a strategy found an acceptable solution. Even though a 10% improvement was the goal, to allow for random error, we defined an acceptable solution as one that had at least 9% improvement. The second metric, called *winning*, is the number of times a strategy found an acceptable solution and moved a shorter distance than all other strategies that also found an acceptable solution. The third metric, called *percent distance from winner*, was measured for each strategy at each starting point and is the percentage difference between the distance moved by a strategy that found an acceptable solution and the distance moved by the winning strategy from that starting point.

As an example, consider Table 1, which shows the raw data for all of the design strategies at a single starting point on a surface of type 10. The response at the starting point for this response surface function was $u = 44.16$, so an acceptable solution is one with a response of at least 9% higher or 48.14. For each strategy, the true response at the solution point for each strategy is shown in the first column. The second column shows the distance from the starting point to the solution point for each strategy.

Table A.1.

Factor	Levels 0	1	2	3	4	5	6
Main effect terms (S)	Just linear	All linear, some quadratic	All linear, all quadratic	All linear, some quadratic & cubic	Low probability of all up to sixth order	Medium probability of all up to sixth order	All up to sixth order
Interaction terms (T)	None	Some two-factor	All two-factor	All two-factor, some three-factor	Low probability of two- three-factor	Medium probability of two- & three-factor	All two- & three-factor
Number of potentially active factors (A)	2–3	4–5	6–7	8–9	10–11	12–13	14–15
Flatness (r)	1	1.33	1.67	2	2.33	2.67	3
Non-fol/ Fol Ratio for S (Rs)	0	0	0.5	0.5	0.5	1	1
Non-fol/ Fol Ratio for T (Rt)	0	0	0.5	0.5	0.5	1	1
Noise (N)	0	0.25	0.5	0.75	1.0	1.25	1.5

Notes: for all the rows in the S matrix, Rs is the ratio of the probability in the second column of the S-matrix to the probability provided the first column. Rt is similarly defined for the T-matrix. N is the variance of the normally distributed error term.

For this starting point, only the FF–CC and FF–SC strategies were successful at finding acceptable solutions. The winner for this simulation was the FF–CC strategy since it found an acceptable solution closest to the starting point. For the *percent distance from the winner*, the FF–SC strategy is

$$[(35.48 - 29.97)/29.97] \times 100\% = 18.39\%$$

from the winner. The *percent distance from the winner* for the FF–CC strategy is zero. For the other strategies, the percent distance from the winner is not calculated because they did not find an acceptable solution.

One interesting outcome was the difference that an error component made on the results. According to the 49 run orthogonal array, the first seven types of surfaces (sets of surface characteristics) had zero variance on their error component, so responses from these surfaces were observed without error. In this case, the OFAAT–BB strategy often found the highest number of acceptable solutions and changed the solution the least to find them. We conjecture that this happens because the OFAAT–BB strategy allows for less model bias, which is the only source of error in the fitted surface for these surfaces with zero error variance. However, this strategy performed poorly in the presence of error, where the FF–CC strategy had much better results on all three metrics. Because of this discrepancy and our feeling that researchers seldom encounter a case where there is no error observed on the responses, we choose to concentrate our analysis on the 42 surface types that were observed with error.

The results for the first metric, success, are displayed in Figure 1. The x-axis on the graph has an entry for each of the 42 surface types that included noise. The x-axis does not display any scale since the surface types were ordered according to the number of times that the FF–CC strategy found an acceptable solution. This ordering of the surface types makes the graph much easier to read. The y-axis shows the number of times out of 125 simulations on that surface type where each strategy was successful at finding at least a 9% improvement. The graph shows that all strategies except the FF–CC strategy, are indistinguishable in terms of the number of successes. The FF–CC strategy dominates all of the other strategies for every surface type where at least one strategy was able to achieve 10 successes out of 125. On the surface types where there are less than 10 successes, the surfaces were very flat relative to the error and thus no strategy was able to perform well.

Figure 2 shows the same graph for the second metric, number of winners, for each strategy and the same conclusions can be drawn. Namely, that the FF–CC strategy dominates all of the other strategies on surfaces that are not extremely flat. Also, the fact that all of the other strategies are indistinguishable in terms of performance for this metric. McDaniel [9] provides details that confirm these conclusions at a 95% confidence level using the method for comparisons with the best suggested by Hsu [10].

Table B.1.

Surface type	N	S	T	A	r	Rs	Rt
1	0	0	0	0	0	0	0
2	0	1	6	6	6	6	6
3	0	2	5	5	5	5	5
4	0	3	4	4	4	4	4
5	0	4	3	3	3	3	3
6	0	5	2	2	2	2	2
7	0	6	1	1	1	1	1
8	1	0	6	5	4	3	2
9	1	1	5	4	3	2	1
10	1	2	4	3	2	1	0
11	1	3	3	2	1	0	6
12	1	4	2	1	0	6	5
13	1	5	1	0	6	5	4
14	1	6	0	6	5	4	3
15	2	0	5	3	1	6	4
16	2	1	4	2	0	5	3
17	2	2	3	1	6	4	2
18	2	3	2	0	5	3	1
19	2	4	1	6	4	2	0
20	2	5	0	5	3	1	6
21	2	6	6	4	2	0	5
22	3	0	4	1	5	2	6
23	3	1	3	0	4	1	5
24	3	2	2	6	3	0	4
25	3	3	1	5	2	6	3
26	3	4	0	4	1	5	2
27	3	5	6	3	0	4	1
28	3	6	5	2	6	3	0
29	4	0	3	6	2	5	1
30	4	1	2	5	1	4	0
31	4	2	1	4	0	3	6
32	4	3	0	3	6	2	5
33	4	4	6	2	5	1	4
34	4	5	5	1	4	0	3
35	4	6	4	0	3	6	2
36	5	0	2	4	6	1	3
37	5	1	1	3	5	0	2
38	5	2	0	2	4	6	1
39	5	3	6	1	3	5	0
40	5	4	5	0	2	4	6
41	5	5	4	6	1	3	5
42	5	6	3	5	0	2	4
43	6	0	1	2	3	4	5
44	6	1	0	1	2	3	4
45	6	2	6	0	1	2	3
46	6	3	5	6	0	1	2
47	6	4	4	5	6	0	1
48	6	5	3	4	5	6	0
49	6	6	2	3	4	5	6

The third metric, the percent distance from winner, is displayed in Table 2. Because the graphical display was confusing to read, the results are shown in tabular form only. Here, the FF–CC strategy had the lowest mean, significantly lower (at 95% confidence) than either the FF–BB strategy and the OFAAT–CC strategy which also did reasonably well on this metric. In the table, N refers to the number of surface types (out of 42) on which the strategy was able to find at least one acceptable solution (out of 125 simulations). On surface types where no acceptable solutions were found by a given strategy, the average percent distance from the winner could not be calculated.

The most striking result of the research conducted here was the superiority of the FF–CC strategy for almost all of the 42 design points. In both the number of times it found an acceptable solution and the number of times it found that solution by making the smallest change from the starting point, the FF–CC strategy was best in almost every case. Additionally, the FF–CC strategy had the smallest average difference in distance from the best. When there is no error in the response the OFAAT–BB strategy often worked better than the FF–CC strategy, but when even the smallest amount of error was present it performed substantially worse. Our recommendation is to use a traditional response surface methodology approach similar to the FF-CC strategy when confronted with the small factor change problem.

APPENDIX A: FACTORS AND LEVELS FOR THE ORTHOGONAL ARRAY

See Table A.1.

APPENDIX B: ORTHOGINAL ARRAY

See Table B.1.

ACKNOWLEDGEMENTS

The authors would like to thank Dan Meyer, Liz Schiferl, Barry Nelson, Ajit Tamhane, Kwok Tsui, and Ted Allen for helpful discussions during this research. The research was supported in part by the Morris E Fine Junior Professorship in Materials and Manufacturing and the Ellis R Ott Scholarship.

REFERENCES

1. Myers RH, Montgomery DC. *Response Surface Methodology: Process and Production Optimization Using Designed Experiments*. John Wiley: New York, 1995.

2. McDaniel WR, Ankenman BE. A response surface test bed. *Quality and Reliability Engineering International* 2000; **16**(5):363–372.
3. NPR (National Public Radio Report by Steve Tripoli). Designed foods. *Morning Edition*, September 24 1997.
4. Hamada M, Wu CFJ. Analysis of designed experiments with complex aliasing. *Journal of Quality Technology* 1991; **24**(3):130–137.
5. Chipman H, Hamada M, Wu CFJ. A Bayesian variable selection approach for analyzing designed experiments with complex aliasing. *Technometrics* 1997; **39**:372–381.
6. Statistical Sciences, Inc. *S-PLUS User's Manual, Version 3.3 for Windows*. Statistical Sciences, Inc.: Seattle, 1995.
7. Bazaraa MS, Sherali HD, Shetty CM. *Nonlinear Programming: Theory and Algorithms*. John Wiley: New York, 1993.
8. Mee R. Follow-up designs that make the most of an initial screening experiment. *41st Annual Fall Technical Conference*. American Society for Quality: Baltimore, MD, 1997.
9. McDaniel WR. Using simulation to compare experimental design strategies for quality improvement with minimal changes to factor levels. *Ph.D. Thesis*. Northwestern University, Department of Industrial Engineering and Management Sciences, 1999.
10. Hsu JC. *Multiple Comparisons: Theory and Methods*. Chapman and Hall: New York, 1996.

Authors' biographies:

William McDaniel received his PhD in Industrial Engineering/Management Science from Northwestern University in 1999, where his research interests included response surface methodology, design of experiments, and simulation. He is currently working as an Operations Research Consultant for ZS Associates in Evanston, IL.

Bruce Ankenman received his PhD in Industrial Engineering from the University of Wisconsin-Madison in 1995. He is now an Assistant Professor in the Department of Industrial Engineering and Management Sciences at Northwestern University. Research interests include design of experiments, robust design, and experiments involving variance components and dispersion effects.

