

COLUMNWISE CONSTRUCTION OF RESPONSE SURFACE DESIGNS

William D. Heavlin and G. Paul Finnegan

Advanced Micro Devices

Abstract: All commonly used, general purpose algorithms for constructing experimental designs work design point by design point-rowwise. We introduce an algorithm that works columnwise, that is, factor by factor. In common with other algorithms, ours requires an optimality criterion with respect to a specified model. Among its advantages are its ability to accommodate a priori notions of symmetry and balance, to adapt experiments with sequentially processed factors, and to incorporate goodness-of-fit considerations. Through a series of problems, we explore the properties and utility of this approach, comparing its solutions to those of other design algorithms in terms of D-optimality, design yield, and level balance.

Key words and phrases: Balance, covariates, missing values, optimal design, semiconductor industry.

1. Introduction

Mainstream experimental design practice centers on factorial and fractional factorial designs and orthogonal arrays. Box and Hunter (1961) introduce two-level fractional factorials for industrial applications, and Box, Hunter and Hunter (1978), chapter 12 provide classical instruction on the confounding relationships incumbent in such designs. Franklin (1985) provides extensions to fractional factorial experiments of prime powers. To construct two-level fractional factorials when only a subset of interactions is of interest, Taguchi (1987), volume I, chapter 7, Kacker and Tsui (1990), and Wu and Chen (1992) adapt linear graphs as teaching aids. Sun and Wu (1994) extend these methods to three level designs. Characteristic of the interaction graph approach is its restriction to orthogonal arrays, trial-and-error flavor and application in stages, each step adding one more factor (column) to a tentative solution.

General algorithms for constructing experimental designs operate differently, using optimality criteria with respect to particular models in order to select points (rows) from the design space. Silvey (1980), chapter 2 and Atkinson and Donev (1993), chapter 10 review various criteria for design optimality. Of these, the most common is the easy-to-compute determinant of the coefficients' variance-covariance matrix, D-optimality, which is asymptotically equivalent to minimum prediction error. Proposals by Van Schalkwyk (1971), Mitchell and Miller (1970),

Wynn (1970), Mitchell (1974a, 1974b), and Galil and Kiefer (1980) all present criterion-optimizing algorithms that act by alternately adding and deleting design points to and from a tentative solution. Fedorov (1972), Cook and Nachtsheim (1980), Johnson and Nachtsheim (1983), and Atkinson and Donev (1989) describe various exchange algorithms, in which current design points are replaced by other candidate points. Algorithmic approaches incorporating covariates or blocks are particularly relevant to the present work. Key references are Harville (1974), Jones (1976), Cook and Nachtsheim (1989), Nachtsheim (1989), and Atkinson and Donev (1989), discussed further in section 4. A succinct survey of optimal design algorithms can be found in Atkinson and Donev (1993), chapter 15.

The advantages of such algorithms include their flexibility with regard to the total number of observation points, the form of the model, the size of blocks, and constraints in the design space. These advantages result in simplified user interfaces, which in turn make such algorithms good supplements to designs based on orthogonal arrays. Of the algorithmic approach, Box and Draper (1987), chapter 14 are critics, first, for its inflexibility in accommodating subject matter expertise, secondly, for its overdependence on single models, thirdly, for the poor handling of model goodness-of-fit issues (e.g. centerpoints, rotatability, and resolution IV concepts).

In this paper, we present an experimental design method—algorithm DR, below—that combines elements of the orthogonal array and algorithmic approaches. In common with the algorithmic approach, our method requires an optimality criterion and a specified model. In common with other orthogonal array methods, ours invites user-specified subarrays whose symmetry and balance characteristics are preserved. As a side benefit, our approach gives new flexibility for assessing model goodness-of-fit.

We classify our algorithm by the term “columnwise”. This is to contrast it to the algorithmic approaches listed above, which operate rowwise, adding, deleting, and exchanging points in the design space, that is, rows of the design matrix. The term “columnwise” is chosen also to establish a certain sympathy our algorithm shares with the constructive interaction graph methods of Kacker and Tsui (1990), Wu and Chen (1992), and Sun and Wu (1994). These orthogonal array-based methods also construct solutions by joining one or more factors (columns) to a partial solution. Our use of the term is consistent with that made by Li and Wu (1997), who apply similar ideas to construct supersaturated designs, and Li (1995), to two-level designs. In addition, Park (1994) proposes a columnwise construction of Latin hypercubes for computer experiments.

Our algorithm seems a natural solution to certain experimental design problems, and these we present in section 2. From these problems, we abstract pertinent data structures in section 3, and in section 4 we describe four algorithms—three rowwise, one columnwise—that we use to solve them. The details to the

solutions are described in section 5, but we reserve most commentary until section 6.

2. Examples

In this section, we introduce selected design problems from the semiconductor industry. Known as integration experiments, the problems are intrinsically interesting, and not easily solved by existing techniques.

Integration experiments are to be distinguished from experimental characterizations of the individual processing steps, typical of the work performed by SEMATECH and other equipment-oriented development efforts. For single-step characterizations, typical responses involve physical dimensions and deposited particles, the factors comprise process recipe values, and mainstream experimental methods work satisfactorily. In contrast, integration experiments involve batches of material processed over multiple fabrication steps. Responses are typically electronic properties of transistors, resistors, capacitors, circuits, and so on, quantities not measurable until the end of processing, while the factors consist of process recipe settings, with two to three levels for each of several steps.

A defining property of integration experiments is the time order in which the experimental material—silicon wafers—is processed. At a given experimental step, the batch of wafers is divided into two or three groups corresponding to the levels of the factor for that step; after all wafers are processed at their assigned levels, the batch is combined again for subsequent processing. At the next step at which multiple levels are investigated, the wafer batch is again divided into groups according to the factor levels associated with that step, the wafers processed, the batch recombined and sent on. The time of processing thus parallels the *column order* or the factor sequence of the experimental plan, as opposed to the row order encountered in conventional, single-step characterizations. (Secondarily, there are interesting time-order and batching effects—split plot issues, randomization, and so on—implied by the processing orders within each step. These are beyond the scope of the present paper.)

2.1. Problem WLR: Wafer loss and repair

Table 1a shows the experimental design planned for one such integration experiment. The 24 rows denote the original layout for each of the silicon wafers of the 24-wafer batch. There are four factors: *p*-well implant dose, threshold voltage adjustment implant dose, and *n*-channel lightly doped drain dose and energy.

After the second factor was completed but before the third and fourth ones are, three wafers are broken; the 24-wafer design is now a 21-wafer design. We can at least contemplate that the assigned levels for factors three and four could

be modified to some benefit. The two solutions in Tables 1b and 1c are described in section 5.1.

Table 1. For problem WLR, (a) the original design before experimental units wafers 2, 12, and 18 were lost, the solution for the subarray X (b) by the columnwise algorithm DR, and (c) by the rowwise algorithm D2.

wafer	W :		X : NLDD		X : NLDD	X : NLDD
	p -well	V_{ta}	dose,	energy		
01	-	-	0	+	-	+
02 lost	+	0	-	-		
03	-	+	-	+	+	+
04	0	-	-	0	+	-
05	-	0	-	0	+	+
06	+	+	+	+	0	+
07	+	-	-	-	+	0
08	+	-	0	-	-	0
09	+	-	+	+	-	+
10	-	0	+	-	-	-
11	+	+	+	+	-	-
12 lost	-	-	-	+		
13	+	0	0	-	-	+
14	-	+	0	+	0	+
15	+	-	-	+	0	-
16	-	+	0	0	-	0
17	-	0	+	0	0	0
18 lost	0	-	-	0		
U:						
19	0	+	+	-		
20	0	+	+	-		
21	0	+	+	-		
92	0	0	0	0		
23	0	0	0	0		
24	0	0	0	0		

(a)

X : NLDD	X : NLDD
dose,	energy
-	+
+	+
+	-
+	+
0	+
+	0
-	0
-	+
-	-
-	-
-	-
-	+
0	+
0	-
-	0
0	0

(b)

X : NLDD	X : NLDD
dose,	energy
-	+
+	0
+	-
+	+
+	+
-	-
-	+
+	+
-	-
-	-
-	-
+	-
-	+
0	-
0	-
+	-

(c)

Wafer breakage is not the only reason this kind of problem might arise. Heavlin and Finnegan (1993) give an example where two previous manufacturing steps introduce covariates; no longer identical, and in spite of likely interactions, the wafers must still be assigned to planned experimental conditions, even though the wafers are not homogeneous. In addition, the natural flow of material through

the factory implies a pipeline of experiments under way. As information from one experiment becomes known, plans for other, partially completed experiments can and should change. Further, there is a delay between designing an experiment and its execution; during that interval experimenters have new ideas. Some of these come from deeper thinking on the part of the experimenter, others are the consequences of encountering unexpected phenomena in other batches. Certain new ideas are not destructive or non-destructive with any certainty, but they are high risk. Finally, occasionally a factor is misprocessed—the wafers intended for certain levels receive others instead. For all these reasons, one may wish to modify the original experiment, taking into account the factors already executed, the experimenter’s current knowledge, and the development effort’s foremost priorities.

2.2. Problem NPC: Simultaneous n -channel, p -channel, and circuit models

The most common technology in use for fabricating contemporary integrated circuits is CMOS (which stands for complementary metal oxide semiconductor). CMOS circuits combine two kinds of transistors, p -channel and n -channel, which carry signals with positive and negative charges, respectively. The p - and n -channel transistors are spatially separate, and, with a few exceptions, they are defined by separate, but analogous, manufacturing steps.

A common practice in designing CMOS characterization experiments is to confound certain factors (e.g. the n -channel specific ones) with certain others (e.g. the p -channel specific ones). For example, the $(-, 0, +)$ -levels of the p -channel transistor’s n -well implant dose are identical to the $(-, 0, +)$ -levels of the n -channel transistor’s p -well implant dose. For work early in a technology’s development, such confounding is acceptable, since the focus is to optimize and stabilize each type of transistor. For later work, however, this practice has unattractive consequences. Responses measured of the CMOS circuit as a whole, speed and power consumption for instance, result from n - and p -channel devices working together. For these circuit-oriented studies, confounded designs are no longer acceptable.

Table 2 lists the eight factors and their relation to the three response categories, p -channel transistors, n -channel transistors, and circuit properties. The two transistor types share only one factor, the threshold voltage adjustment implant dose, but the factors affecting the circuit responses are nearly a union of those of the two transistors. (Indeed, this is literally true; however, for physical reasons, the n -well and p -well doses are thought to have only a second-order effect on circuit behavior.) Consistent with mainstream experimental design practice, and feasible with four or five factors and 24 wafers, the designs for the p - and

n -channel responses are desired to fit the full quadratic models, including all second-order interactions. For the circuit model, with six to eight factors, the experimenter recognizes that some interactions need to be given priority over others and has listed eight physically plausible interactions.

Table 2. The relationship among the eight factors for problem NPC across the three response sets. The circuit model also includes the following eight interactions: (i) $V_{ta} \times \text{PLDD}_1$ dose, (ii) $V_{ta} \times \text{PLDD}_1$ energy, (iii) PLDD_1 dose \times PLDD_2 dose, (iv) PLDD_1 energy \times PLDD_2 dose, (v) $V_{ta} \times \text{NLDD}$ dose, (vi) $V_{ta} \times \text{NLDD}$ energy, (vii) PLDD_1 dose \times PLDD_2 energy, (viii) NLDD dose \times NLDD energy.

response	n -well	V_{ta}	PLDD_1 dose	PLDD_1 energy	PLDD_2	p -well	NLDD dose	NLDD energy	model
p -channel	•	•	•	•	•				full quadratic
circuit		•	•	•	•		•	•	8 interactions
n -channel		•				•	•	•	full quadratic

Were the circuit-specific responses not of interest, the solution is straightforward: Obtain a D-optimal design for the p -channel factors and its five-factor response surface model. This determines also the levels of n -channel's threshold voltage adjustment dose. Then associate n -channel's p -well factor with p -channel's n -well, and its NLDD dose and energy levels with those of PLDD_1 . One referee correctly notes that such a four-factor subset of a D-optimal five-factor design is not necessarily D-optimal with respect to the corresponding four-factor model subset. Of course, the more pressing consideration is that such an approach solution does not allow the circuit model to be estimated. Solutions for all three response sets are in section, 5.2.

2.3. Problem 24W: 24-row designs

Underlying problems WLR and NPC is a batch size of 24 wafers. At this writing, this is the size for most advanced semiconductor manufacturing lines, and designs of this size are of interest to our clients and those of many of our colleagues. Also, 24 is an experiment size of some intrinsic interest: Five-factor-three-level full quadratic models (with 21 terms) and six-factor-two-level resolution V models (with 22 terms) both fit within an experiment of 24 units. In contrast to three- and four-factor designs, five- and six-factor designs are large enough to not be easily visualized. Experiments with about five or six factors are also of considerable interest for integration experiments. Finally, 24 is not a power of two, yet rich in small prime numbers, giving plenty of opportunities for symmetry. 24 is larger than 16, a well characterized world in experimental designs, yet smaller than 32, a domain of experiments outside mainstream practice.

For all these reasons, it is attractive to develop a standard set of experimental designs with five or six factors and 24 experimental units, for all models with second-order interaction and quadratic terms. Let us introduce the self-explanatory notation $2^k 3^{F-k}(n)$ to denote an experiment with n experimental units, F factors, of which k are at two levels and $F - k$ are at three levels. Note that a response surface model for a $2^k 3^{F-k}$ design has $1 + F(F + 1)/2 + (F - k)$ terms: 1 (constant) + F (linear) + $F(F - 1)/2$ (interactions) + $(F - k)$ (quadratic). The standard set of designs we seek is therefore $\{2^k 3^{5-k}(24) : k = 0, \dots, 5\} \cup \{2^k 3^{6-k}(24) : k = 4, 5, 6\}$, a set which has nine elements and whose models contain 16 to 24 terms. However, as indicated by problem WLR, we have considerable interest in design properties after random loss of some experimental units. Therefore, we confine our attention to $\{2^k 3^{5-k}(24) : k = 0, \dots, 5\} \cup \{2^6(24)\}$, which allows for the loss of at least two units in each case. We describe the solutions in section 5.3.

3. General Problem Statement

In this section, we abstract key elements of problems WLR, NPC, and 24W.

3.1. Matrix structures

Let the design matrix H be partitioned as $\begin{bmatrix} W|X \\ U \end{bmatrix}$, where the matrices W and U are given and X is to be determined. (A typical use of U is to force into the solution one or more centerpoints.) This partitioning is manifestly part of problem WLR, but we acknowledge that its value for problems NPC and 24W is less apparent.

These data structures can be used to represent the covariates and blocking variables, so our three problems resemble those addressed by Harville (1974), Cook and Nachtsheim (1989), Atkinson and Donev (1989), Nguyen (1994), and Radson and Herrin (1995), for instance. An important distinction, developed next, is that our models permit interactions among the factors comprising W and X . In contrast, the latter authors focus on additive models, which they are able to exploit for computational efficiency.

3.2. Models

From H , whose i th row is H_i , we have one or more model matrices M^j , indexed by j and whose i th row is given by the vector-valued function $m_j(H_i)$. For example, when H is a two-column matrix, and m corresponds to a full quadratic model, $m(u_1, u_2)$ returns the row vector $(1, u_1, u_1^2, u_2, u_1 u_2, u_2^2)$. The row-by-row transformation of H into M^j we abbreviate to the expression $M^j = m_j(H)$. Problem NPC presents us with three m -models, while problems WLR and 24W require but one.

3.3. Criteria

In our solutions to problems WLR, NPC, and 24W, we consider three measures, D-optimality, design yield, and level balance, defined here.

(a) *primary criterion*, D-optimality: Our goal is to find the best X , or at least a good one, so that we can estimate by least squares the linear coefficients of M^j . Different criteria for design optimality have been proposed, the most common of which is the larger-is-better criterion called D-optimality: $D(H) = \ln(\det(M^T M))$, when $M = m(H)$ is non-singular; otherwise, $D(H) = -\infty$. Among the advantages of $D(\cdot)$ is the ease with which it is computed. A review of alternative criteria can be found in Silvey (1980), chapter 2 and Atkinson and Donev (1993), chapter 10.

When there are multiple models m_j , the definition of $D(\cdot)$ is naturally generalized as $D(H) = \sum_j \ln(\det(M^{jT} M^j))$, for all $M^j = m_j(H)$ non-singular; $D(H) = -\infty$, otherwise. This multiple model version of D-optimality, sometimes referred to as S-optimality, has been studied by Läuter (1976) and Atkinson and Cox (1974); Cook and Nachtsheim (1982) do so using as criteria linear functionals of $(M^{jT} M^j)^{-1}$.

(b) *secondary criterion*, design yield: Given that exactly w of n experimental units are lost, and lost at random with equal probability, we define design yield (w, m_j) as the probability that the remaining $n - w$ rows of M^j comprise a non-singular design. For example, in the case of Table 1a, 21 wafers remain. Breaking any one of these, we can still support the four-factor quadratic model, and the design yield (1) is 100 percent. If two wafers are broken, however, 78 of the 210 possible pairs give singular designs, and the design yield (2) is $(210 - 78)/210 = 62.9$ percent; for three wafers, 923 of 1330 possible triples are singular, the design yield (3) is 30.6 percent; for four wafers, 11.6 percent. As a function of w , we usually present design yield(w) either graphically as in Figure 2 or numerically as in Table 3.

Table 3. For the three solutions to problem WLR after wafers 2, 12 and 18 are lost, the number of singular designs that would occur assuming random wafer breakage of 1 to 6 more wafers. Design yield is 1 minus the tabled value divided by the first column. For example, for the original design and two wafers, the design yield is $1 - (78/210) = 62.9$ percent.

			original	DR	D2
of	21	1-wafer losses	0	0	0
of	210	2-wafer losses	78	2	1
of	1330	3-wafer losses	923	467	461
of	5985	4-wafer losses	5291	4125	4112
of	20349	5-wafer losses	19718	18264	18252
of	54264	6-wafer losses	53993	53246	53243

Herzberg and Andrews (1976) and Andrews and Herzberg (1979) formulate design yield as follows: Let $\{d_i : i = 1, \dots, n\}$ denote independent, identically distributed Bernoulli random variables; d_i takes value 0 when wafer i is missing and value 1 when wafer i is present. Define $D = \text{diag}(d_i)$. These authors consider the criteria $\Pr\{\det(M^T DM) \neq 0\}$ and $E\{n^{-1} \det(M^T DM)^{1/p}\}$, where p denotes the number of columns of M . Using this notation, design yield(w) is the conditional probability $\Pr\{\det(M^T DM) \neq 0 \mid \sum d_i = n - w\}$, which we note is free of the Bernoulli parameter $\Pr\{d_i = 1\}$.

(c) *constraint*, level balance: Let $H^{(f)}$ denote column f of the design matrix H . We say that a column or factor f is level balanced (or has balanced levels) if the unique scalar elements in column $H^{(f)}$ occur with equal frequency. The relevance of level balance is discussed further in section 5.

In applications, based on a random start, the algorithm DR maximizes criterion (a) to identify a candidate solution. The algorithm is repeated three to ten times. These solutions are all evaluated by criterion (b), and the “best” solution is selected with both criteria in mind. We take this approach because criterion (b) is both too slow to make it the primary objective function for algorithm DR, yet also too important for our clients to disregard entirely.

4. Algorithms

In this section, we introduce four algorithms for solving the problems WLR, NPC, and 24W. All algorithms seek to optimize criterion $D(\cdot)$, although the fourth, DR, does so over a restricted domain.

4.1. Algorithm JN

This is the implementation of the algorithm of Johnson and Nachtsheim (1983) with $k = 3$. Due to its computational efficiency, JN is the most commonly available D-optimal algorithm. In particular, we utilize the implementation found in RS/Discover, described further by Bolt, Beranek, and Newman (1992), appendix 3. According to Johnson and Nachtsheim (1983), and confirmed by Li and Wu (1997), algorithm JN is equivalent to the Fedorov algorithm with $k = n$.

4.2. Algorithm D1

D1 is the Fedorov algorithm, replacing one design point at a time to maximize the D-optimality criterion. It has been suitably modified to leave W intact.

Let $q = \dim(X)$ and let $\mathcal{X} = \{-1, 0, +1\}^q$. Given some starting design X_0 , we update it with X_1 , with the following properties: (a) X_1 is the same as X_0 except in row i^* . (b) row i^* of X_1 is replaced with some q -dimensional point x^* in \mathcal{X} . (c) $D(\cdot)$ is larger for (i^*, x^*) than for any other choice (i, x) . This algorithm is repeated until no further improvement can be made.

4.3. Algorithm D2

D2 is similar to algorithm D1 except that it changes up to two points at a time: Given some starting design X_0 , we update it with X_1 with the following properties: (a) X_1 is the same as X_0 except in row i^* and perhaps row j^* . (b) row i^* of X_1 is replaced with some q -dimensional point x_i^* in \mathcal{X} ; row j^* of X_1 by x_j^* . (c) $D(\cdot)$ is larger for (i^*, j^*, x_i^*, x_j^*) than for any other choice (i, j, x_1, x_2) . This algorithm is repeated until no further improvement can be made.

4.4. Algorithm DR

DR stands for “design repair.” Let X denote some matrix selected by the experimenter, and let W_i and X_j denote the i th and j th rows of W and X respectively. Let π denote a permutation of the row indices $(1, \dots, n)$, where n here denotes the number of rows of W and X . For any given permutation π , we associate row i of W with row $\pi(i)$ of X , so each row of W is associated with one and only one row of X . Criterion $D(\cdot)$ is used to guide choosing which permutation to use.

Let Π_{all} denote the set of all permutations, and let $\Pi_1(\pi)$ denote the subset of permutations formed by transposing at most one pair of indices of π ; $\Pi_1(\pi)$ consists of $1+n(n-1)/2$ elements. For example, $\Pi_1((1, 2, 3)) = \{(1, 2, 3), (2, 1, 3), (3, 2, 1), (1, 3, 2)\}$ but not the two cyclic permutations $(2, 3, 1)$ and $(3, 1, 2)$.

Algorithm DR has two steps, a *random* starting point from Π_{all} (called R-step) and a deterministic search over the *exchanges* or transpositions from Π_1 (called E-step). *R-step*: π is selected at random from Π_{all} for n_R iterations and that with the largest D -value is selected. Denote this permutation π_R . *E-step*: Let $\pi_0 = \pi_R$. The $k+1$ st step of E-step takes that permutation π from $\Pi_1(\pi_k)$ that maximizes $D(\cdot)$. E-step continues until the stopping condition $\pi_{k+1} = \pi_k$ is obtained. (A referee notes that E-step corresponds precisely to Harville’s (1974) interchange step.) *Specification of n_R* : The derivation deferred to the appendix, we take n_R from the following relation: $\log_{10}(n_R) \doteq -0.70850 + 2.12105 \log_{10}(n)$. Except for the initial choice of X , this specifies the DR algorithm completely.

In practice, algorithm DR’s ultimate solution depends on the starting point found by R-step, and the different solutions often achieve different values of $D(\cdot)$. In contrast, the rowwise algorithms, particularly JN, are more repeatable. As a matter of practice, the authors run algorithm DR somewhere between three and ten times, and the final solution is selected by reviewing both the D -values achieved and their design yields. The DR solutions of this paper are selected from the best of ten alternatives.

Because DR is run up to 10 times, a referee raised the issue of whether the comparison of DR to the other algorithms is fair. We believe it is, for the following

reasons: (a) On the criterion of $D(\cdot)$, the initial design is always algorithm DR's solution. All the rowwise algorithms easily improve upon it by dropping a few points with 0-levels in favor of those with values ± 1 . (b) On the criterion of design yield, we emphasize the average design yield across the 10 DR solutions, not the best of 10.

Harville (1974) first considered the problem of assigning covariates into a design with qualitative variables. His algorithm consists of two phases. The first ("exchange") replaces the treatment to a given experimental unit with a different treatment. The second ("interchange", equivalent to E-step above) switches the assignment of treatments between two experimental units. Jones (1976) develops a blocking algorithm for a weighted trace; its flexibility for response surface models is less apparent. Nachtsheim (1989) develops the asymptotic theory for covariate assignment in experimental designs. Atkinson and Donev (1989) formulate the KL-exchange algorithm for blocking response surface models. To the same problem, Cook and Nachtsheim (1989) adapt Galil and Kiefer's (1980) proposal for the starting design, and Harville's algorithm for the refinement. All these approaches gain computational efficiency by assuming the absence of interactions between covariates and blocks on one hand (W -structures) and experimental conditions (X -structures) on the other.

A few comments are in order. First, none of these algorithms for designing experiments carries any guarantee of being truly optimum. Second, all are dependent on some starting layout for X . We find it convenient to use the solution from DR as starting points for D1 and D2. Third, algorithm DR mandates more care in choosing its initial X than do JN, D1, and D2, for its rows are never replaced, only permuted. Fourth, algorithms JN, D1, and D2 all comprise rowwise algorithms, making selections and exchanges of design points from the space \mathcal{X} ; this same property allows constraints in the design space. DR, depending as it does on the experimenter's initial choice of X , we describe as a columnwise algorithm. In the solutions discussed next, we treat algorithms JN, D1, and D2 as representatives of the class of rowwise algorithms, select whichever is the most suitable, and compare its performance to DR.

Finally, note that none of the four algorithms makes use of design yield as an optimization criterion. Andrews and Herzberg (1979) develop a recursive relationship for calculating the determinant with one row missing given the complete determinant and design matrix; Akhtar and Prescott (1986) apply this recursion to determine designs optimizing $E\{\det(M^T DM)^\lambda\}$, for $0 < \lambda < 1$. For our problems, this criterion appears impractical: Problem WLR, our smallest, has a model of 15 terms and 21 experimental units; a single criterion evaluation involves potentially 82,160 combinations.

5. Solutions to Problems

5.1. Solutions to WLR

Algorithm DR's solution is displayed in Table 1b. Implicit in Table 1b is the experimenter's choice of rows for DR to permute, and close inspection reveals that enforcing level balance underlies that choice.

Since the \mathcal{X} -space consists only of 9 points, we can use algorithm D2 to good effect. Its solution is presented in Table 1c; as a starting design algorithm D2 used Table 1b. By construction, therefore, D2's has a superior D -value to DR, as can be seen in Table 4.

Table 4. D -optimality criterion values achieved different algorithms for the three design problems. Higher values are better. Symmetry designs refer to the $3/4$, $1/2$ and $1/3$ fractional factorials.

	original	columnwise	rowwise	symmetry
problem WLR	24.76	27.75	31.17	
problem NPC				
n -channel	33.91	34.24	35.64	
circuit	$-\infty$	44.17	45.58	
p -channel	45.67	45.67	48.88	
problem 24W				
2^6		65.92	68.01	
3^5		45.67	49.23	
$2^1 3^4$		46.66	50.79	
$2^2 3^3$		48.06	50.81	
$2^3 3^2$		48.51	50.54	49.27
$2^4 3^1$		49.84	50.25	50.14
2^5		49.91	49.91	49.91

Observe that D2's solution is not at all balanced: The third factor has 6 -1 's, 5 0 's, and 10 $+1$'s, the fourth 11, 4, and 6 -1 's, 0 's and $+1$'s, respectively.

A comparison of design yields is in Table 3, where the solutions perform comparably. Both improve over Table 1a by the equivalent of around one experimental unit, an improvement typical of the authors' experience with an experiment of about this size. Note that when two wafers are lost, D2's design yield is 0.995, versus DR's value of 0.99.

5.2. Solutions to NPC

For NPC, the \mathcal{X} -space consists of $(-1, 0, +1)^8$, which has 6561 distinct points. For computational reasons, the rowwise solution uses algorithm D1 and optimizes

the sum of D -values of the three models presented in Table 2. The solution from algorithm DR is used as a starting design.

Algorithm DR's solution to problem NPC is sequential; at each step, only one model is used. In step 1, the $3^5(24)$ design from problem 24W is assigned to the p -channel factors.

In step 2, to complete the circuit model, the p -channel design is augmented with the two factors of n -channel lightly doped drain (NLDD) dose and energy. For calculating D -values, the model consists of constant, linear, and quadratic terms of all seven factors, and the interactions as listed in Table 2—in effect, a superset of Table 2's specification. The X -matrix is almost a three-times replicated 3^2 factorial; enforcing level balance, only the points $(-1, 0)$, $(0, +1)$, and $(+1, -1)$ are deleted. To ensure that the solution includes a centerpoint, the matrix U consists of a row of zeros.

In step 3, to complete the n -channel model, the design needs only to be augmented by the factor corresponding to p -well dose implant. The model consists again of a superset—constant, linear and quadratic terms for all eight factors, and the n -channel model's interactions as listed in Table 2.

The latter two stages are illustrated in Figure 1. The three matrix portions corresponding to W , X , and U are implied by the left-hand panels, with row and column dimensions in their respective lower right-hand corners. The models are displayed in the panels on the right. In the model glyphs, the $(1, 1)$ -cell is the constant term; its inclusion is implied when the cell holds a "C". The first row indicates the linear terms of the model; their inclusion is implied by an "L". The main diagonal corresponds to the quadratic terms, which, when included, are denoted by a "Q". Any other cell represents the interaction of the factors of that row and column, and the term is active when an "I" occupies it.

Let us denote the construction sequence implemented— p -channel, then circuit, lastly n -channel—as PCN. Though admittedly problem dependent, PCN seems the natural one to the authors. First, the solution to 24W gives a p -channel design associated with a full quadratic model as a starting point. Second, PCN first adds two factors (for "C"), then only one factor (for "N"), modest DR applications of decreasing ambition. In contrast, the construction sequence NCP awkwardly would add three factors (for "C"), then one (for "P"). PNC and NPC are impossible construction sequences, because p - and n -channel factors form a superset of circuit factors. Construction sequences CPN and CNP are awkward at best. They require four steps, not three, because an initial six-factor circuit design would need to be constructed. Further, the six-factor circuit model is not full quadratic, indeed cannot be for a 24-run design, and therefore seems an odd choice for an initial design.

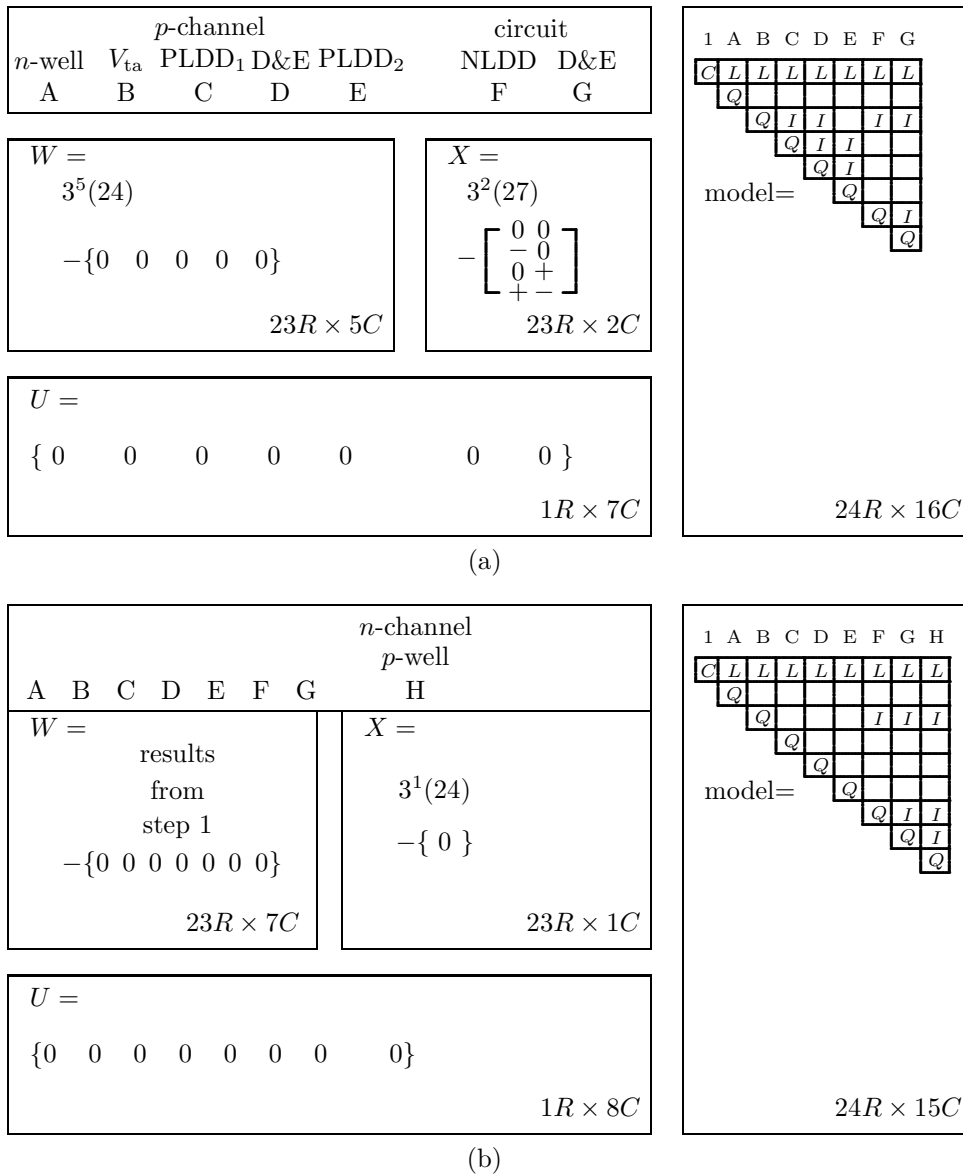


Figure 1. Representation of the two-step solution to problem NPC by the columnwise algorithm DR. (a) In the first step, algorithm DR associates with the *p*-channel design the two additional factors NLDD dose and energy, completing the circuit model. (b) In the second step, algorithm DR associates the result of step one with the factor *p*-well implant dose, completing the *n*-channel model.

Again, as seen in Table 4, algorithm D1 achieves larger *D*-value for all three

models, while DR's solution ensures the levels are balanced. Across the eight factors, the number of 0-levels ranges from 4 to 6, and only for two factors does the number of -1 's equal the $+1$'s. One factor of D1's solution, n -channel lightly doped drain dose, has 11 $+1$'s but only 8 -1 's.

As indicated by Table 5, for both algorithms and all three models the design yields are quite satisfactory. The design yields of the n -channel model distinguishes most between the two algorithms, and shows DR's solution to be more resistant to random observation loss.

Table 5. For the solutions to problem NPC, the number of singular designs assuming random wafer breakage. Values have same interpretation as in Table 3. The rowwise algorithm is D1, the columnwise algorithm is DR, and the original design is $3^5(24)$ from problem 24W.

	original	columnwise	rowwise
<i>p</i> -channel model			
of 24 1-wafer losses	0	0	0
of 276 2-wafer losses	0	0	0
of 2024 3-wafer losses	0	0	0
circuit model			
of 24 1-wafer losses	24	0	0
of 276 2-wafer losses	276	0	0
of 2024 3-wafer losses	2024	0	0
<i>n</i> -channel model			
of 24 1-wafer losses	0	0	0
of 276 2-wafer losses	0	0	0
of 2024 3-wafer losses	0	0	0
of 10626 4-wafer losses	0	0	1
of 42504 5-wafer losses	0	0	23
of 134596 6-wafer losses	6	1	262
of 346104 7-wafer losses	179	52	2119

5.3. Solutions to 24W

Design problems with 24 experimental units, five and six factors, and conventional response surface models, are in the realm of commercial D-optimal software. In our case, we use RS/Discover algorithm, which implements algorithm JN.

Three of these designs can be developed using more traditional fractional factorial methods. John (1971), cited by McLean and Anderson (1984), appendix 4, has formulated a $2^5 3/4$ fractional factorial. This design joins the 16-run 2^{5-1} half-fraction and the 8-run 2^{4-1} half-fraction with one factor fixed. For the $2^4 3^1$ and $2^3 3^2$ there exist 24-run half- and third-fractions, respectively; these are

obtained by coding levels 0, 1 and 2, and taking the fractions whose row sums equal 0 modulo 2, for $2^4 3^1$ and 0 modulo 3, for $2^3 3^2$. As a group, we refer to these designs as “symmetry” designs.

For the DR algorithm, each of the nine designs requires specifying the matrices W , X , and U . Underlying these choices is an interest in symmetry, manifested in the following rules of thumb: (a) Where possible, designate W and X (including U) as full factorials. W and X represent projections into subsets of the factor space. By designating them as full factorials, we are covering these two subspaces as well as possible. (b) Enforce balanced levels. Balanced levels minimize the correlation of effects to a model’s intercept term, facilitate the estimation of the “additive effects” model, and aid in forming useful goodness-of-fit diagnostic plots. (c) Avoid excluding a point with zeros on more than one factor. (d) Avoid including points with zeros on more than one factor. These latter two criteria minimize the concentration of points at only a few radii. As with level balance, the “additive effects” model and goodness-of-fit tests both benefit. (e) Partition the factors so that W has as many factors as possible (and therefore X has as few as possible). For example, to use algorithm DR to solve for a resolution V $2^5(16)$, a natural partition is for W as the 2^4 full factorial and X as a single column of $8 + 1$ s and $8 - 1$ s. The alternative partition into 3 and 2 factors seems less attractive. A detailed review of the implementation for 24W reveals that these rules of thumb are not equally enforced; the latter order indicates something of their relative priority.

Table 4 compares the solutions from algorithm JN and DR with the symmetry designs. For the 2^5 designs, all approaches achieve exactly the same D -value. For all other design problems, the JN solutions, unconstrained by the level balance requirement, achieves uniformly larger D -values, and the symmetry designs for $2^4 3^1$ and $2^3 3^2$ also have marginally better D -values than what is achieved by DR. The disparity in D -values increases as the number of three-level factors increases.

A review of the level balance across the problem sets reinforces this interpretation.

Performance with respect to design yield, however, is more varied. Figure 2 compares the design yield curves for all the five-factor designs and 2^6 . Only for two designs, $2^4 3^1$ and 2^6 , does algorithm JN demonstrate superior design yields over DR’s average performance; for five designs algorithm DR’s best-of-10 solutions show better performance; for 2^6 the results are mixed. All three symmetry design perform poorly.

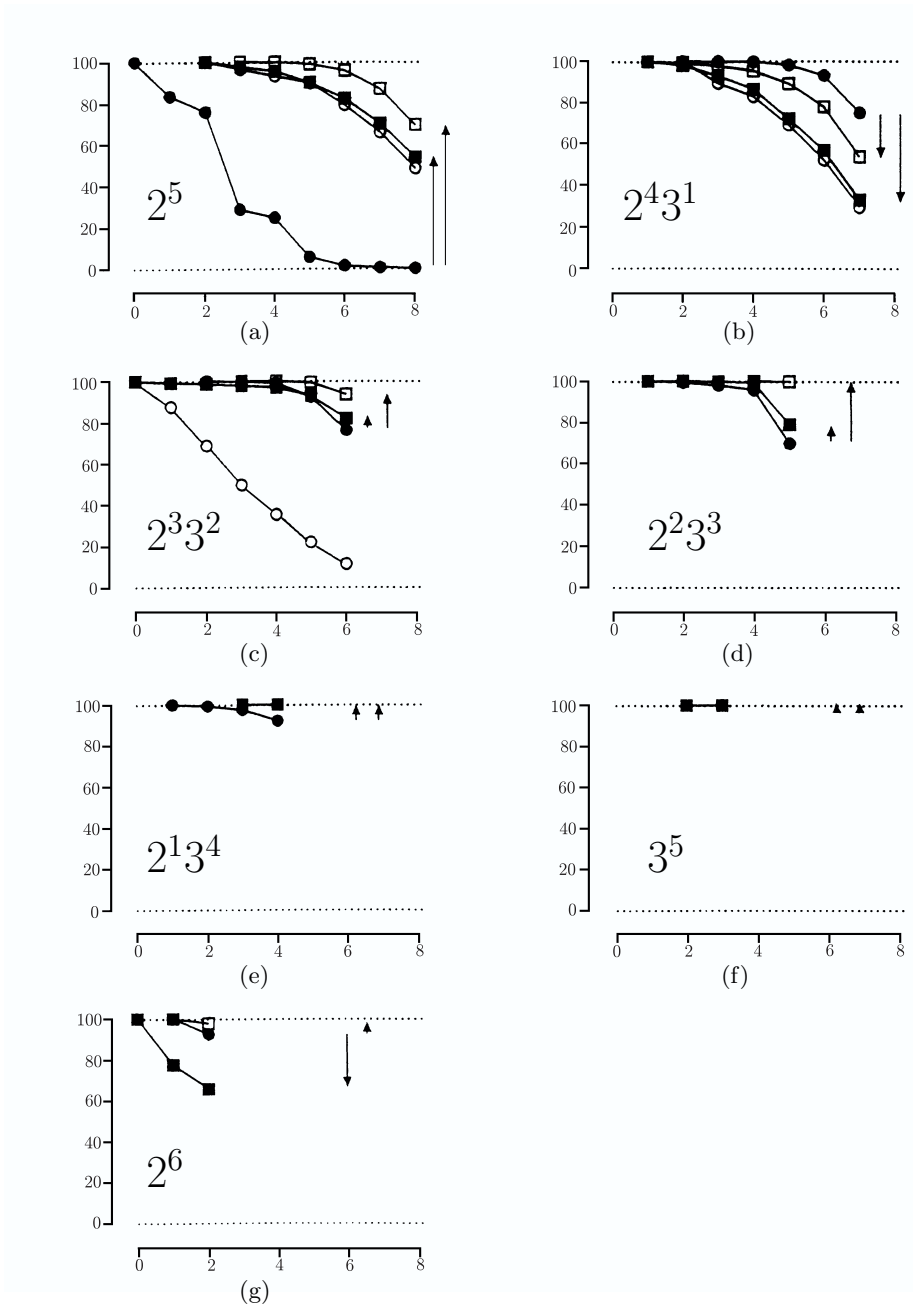


Figure 2. Design yields for the solutions to problem 24W. Symbols: ● denotes the solution from the rowwise algorithm JN, □(■) the best (average) of 10 columnwise solutions from algorithm DR, ○ the 3/4, 1/2 and 1/3 fractional factorials. The arrows compare the magnitude and direction of the rowwise design yield to that of the algorithm DR.

6. Comments

The forte of optimal design theory is solving irregular design issues. Its ability to solve problems WLR, NPC, and 24W we attribute to the judicious choice of data structures: (a) W can consist of previously processed factors, covariates associated with experimental material, and dummy variables indicating block structure, as well as factorial designs to be fractionated algorithmically. (b) U can consist of centerpoints, replicated corner control points, previous designed and complementary blocks, and other conditions that support secondary experimental objectives. (c) The multiple models $\{m_j\}$ capture the experimenter's physical knowledge and primary purposes. With software that supports multiple models, the experimenter is able to specify ambiguous models. (For example, one can represent the resolution IV concept by a series of models, each of which consists of one factor's interactions with all others.) (d) The design points to be, as indicated either by the factorial space \mathcal{X} or the rows of the X -matrix, provide alternative forms for posing design problems. The rowwise algorithms' \mathcal{X} -representation emphasizes the optimality objective and easily accommodates design space constraints; the columnwise algorithm's X -representation gives the experimenter more latitude in specifying solutions that are constrained by symmetry relations, sensitive to model goodness-of-fit, and adapted to less formal considerations. Suppliers of commercial optimal design software are well advised to support data structures (a)-(d).

It is intriguing that algorithm DR can construct a solution to problem NPC through a sequence of one-model-at-a-time steps. This we interpret as further evidence of the strength of the W - U - X data structures, not as a reason to avoid supporting multiple models. One reviewer asks whether DR's success at a sequential construction for NPC is related somehow to the fact that a 24-trial was attempted. Our experience with NPC-type problems is not broad, but it seems likely that problem NPC is near the outer limits of what experimenters in semiconductor development attempt. In addition, for reasons listed in section 2.3, 24-trial experiments seem a natural domain for the application of DR.

Across problems, the rowwise D-optimal algorithms achieve superior values for the D-optimality criterion. Primarily, this is because their solutions are not constrained to have balanced levels. Their disfavor for mid-points (0-levels) is easily explained: Consistent with design practice with two-level designs, the corner vertices give maximal information about both linear and interaction effects, while the 0-levels give new information about the quadratic terms—and there are more of the former than of the latter. The absence of balance between the -1 and $+1$ levels among the rowwise solutions is more problematic, however, and some practitioners are uncomfortable using such solutions for these reasons. The intuition, we believe, is that symmetric designs are more robust to alternative

model specifications, and more supportive of model-free and model-diagnostic data analyses.

In this context, it is noteworthy to observe that the three fractional factorial designs—John’s $3/4\ 2^5$, the half-fraction of 2^43^1 , and the third-fraction of 2^33^2 —show good D -values yet demonstrate relatively poor design yields. Caution is therefore advised in assuming that robustness accrues automatically from design symmetry. From the solutions of problem 24W, the following trends are worth noting, if only speculatively: (1) Level balancing, per se, does not seem to improve design yields. Otherwise, the half-fraction $2^43^1(24)$ and the third-fraction $2^43^1(24)$ would have performed better. (2) Increasing the number of three-level factors improves design yields—albeit the $2^43^1(24)$ designs constitute an important exception. (3) One can improve the design yields of algorithm DR by picking the best of several solutions. In view of (1), it seems the success of DR on the design yield criterion is more easily attributable to its substantial random component (R-step) than to its use of balanced levels. This opens the door to improving the design yields of rowwise algorithms by making their solutions less repeatable.

For all design yield curves presented in this paper, if one design becomes superior to another at one value of w , it remains superior for larger values. (The authors have observed exceptions; this is only a rule of thumb, not a theorem.) This presents the possibility of adapting the proposals of Andrews and Herzberg (1979) and Akhtar and Prescott (1986) to a more computationally feasible, conditional criterion, such as $E\{\det(M^TDM)^\lambda \mid \sum d_i = n - w_0\}$, for $0 < \lambda < 1$ and $w_0 = 1$ or 2 . Were $w_0 = 1$ chosen, the criterion would take at most only n times longer to calculate.

We conclude with some remarks on model goodness-of-fit. Traditional experimental design methods associate goodness-of-fit with resolution IV designs, which designate parts of the model ambiguously. Scientific practitioners tend to associate model goodness-of-fit with the consistency of experimental observations with a physically-based functional form, often made in a single factor’s direction. With the exception of the nonparametric approach of Sacks, Welch, Mitchell, and Wynn (1989), optimal design theory typically neglects goodness-of-fit issues. Algorithm DR offers an approach to accommodating model goodness-of-fit: the user restricts the solutions, forcing in three, four, or even more levels. In so doing, the experimenter has the flexibility to consider model goodness-of-fit directly. The approach of Sacks, Welch, Mitchell, and Wynn (1989) manages goodness-of-fit indirectly through the optimality criterion; algorithm DR does so through user choices, where our experience indicates it belongs. These comments echo the underlying thesis of Cook and Nachtsheim (1989)—that optimal design theory both broadens useful options and further aligns statistical theory to applications.

Acknowledgements

The authors are indebted to Ray Marr for guidance with the combinatorial theory underlying the derivation in the appendix, and to editors and referees for comments that improved the exposition.

Appendix: Derivation of n_R

Let $\Pi_k(\pi)$ denote the set of permutations within k transpositions of the permutation π . Thus $\Pi_0(\pi) = \{\pi\}$ and $\Pi_k(\pi) = \cup\{\Pi_1(\pi') : \pi' \in \Pi_{k-1}(\pi)\}$. For any π , the sets $\Pi_k(\pi)$ are nested: $\Pi_{k-1}(\pi) \subseteq \Pi_k(\pi)$. Note that $(\Pi_{\text{all}} = \Pi_{n-1}(\pi))$ for all permutations π . Define $d(\pi_0, \pi_1)$ as the minimum number of transpositions needed to transform π_0 to π_1 ; $d(\pi_0, \pi_1)$ is the smallest number d such that $\pi_0 \in \Pi_d(\pi_1)$; $d(\pi, \pi) = 0$. $d(\cdot, \cdot)$ is a metric: non-negative, symmetric, and satisfying the triangle inequality.

In the context of algorithm DR, let π^* denote the “best” permutation and let π denote some permutation chosen randomly from Π_{all} with uniform probability. Our heuristic assumes (a) π^* to be unique, and (b) that the maximization of $D(\cdot)$ is similar to the minimization of $d(\cdot, \pi^*)$. Since n_R involves the relative balance of computation time for R- and E-steps, we judge these heuristics adequate for our purpose.

Define the cumulative distribution function $F_n(d)$ to be $\Pr\{d(\pi, \pi^*) \leq d\}$, where π is the random variable. Let $\overline{F}_n(d) = 1 - F_n(d)$. From combinatorial theory, for example, Riordan (1958), chapter 4 $F_n(d) = \sum_{i=1}^d |S_n^{(i)}|/n!$, where $\{S_n^{(i)}\}$ are (signed) Stirling numbers of the first kind. The basic ideas for this result are two: (1) The number of transpositions required to transform one permutation into the identity permutation equals n minus the number of cycles. (2) A cycle of order i factors into $i - 1$ transpositions. (Stirling numbers of the first kind are such that $\Pi_{j=0}^{n-1}(x - j) = \sum_{i=1}^n S_n^{(i)} x^i$. Subject to the boundary conditions $S_n^{(0)} = S_n^{(n)} = 0$, they can be computed recursively: $S_n^{(i)} = S_n^{(i-1)} - (n - 1)S_{n-1}^{(i)}$. In this notation, $(-1)^{n-i} S_n^{(i)}$ is the number of permutations of n objects that have exactly i cycles. See Abramowitz and Stegun (1970), section 24.1.3.)

Let π^* represent the particular permutation maximizing $D(\cdot)$ over Π_{all} . After n_R random permutations, the distribution of that closest in the d -metric is $1 - \overline{F}_n(d)^{n_R}$. Let π_R denote this solution from R -step. Our heuristic is that π_R has a similar distribution, and that E-step moves ever closer in the d -metric. Such ideas allow us to write down the expected number of D -evaluations until convergence:

$$n_R + \binom{n}{2} E(d(\pi_R, \pi) | n_R). \quad (\text{A.1})$$

Here $E(\cdot)$ denotes the expectation with respect to the distribution $1 - \overline{F}_n(d)^{n_R}$. For any non-negative, integer-valued distribution $G(d)$, its expectation is

$\sum_{d=0}^{\infty} \bar{G}_n(d)$. Applied to (A.1), $E(d(\pi_R, \pi)|n_R) = \sum_{d=0}^{n-1} \bar{F}_n(d)^{n_R}$. Because F_n can be calculated, for each choice of n we can minimize (A.1) with respect to n_R ; this minimizing value we denote by n_R^* . The relationship between n and n_R^* is nearly linear on log-log scale. Disregarding a slight concavity, and to excessive precision, a least squares fit on log-log scale through the evaluations at $n = 6, 12, 24, 48, 96$ and 192 gives this approximation: $\log_{10}(n_R^*) \doteq -0.70850 + 2.1205 \log_{10}(n)$. In the region of greatest interest, $n = 12$ or 24 , this expression underestimates n_R^* by 5 or 6 percent.

References

- Abramowitz, M. and Stegun, I. R., editors. (1965). *Handbook of Mathematical Functions, with Formulas, Graphs, and Mathematical Tables*. Dover, New York.
- Akhtar, M. and Prescott, P. (1986). Response-surface designs robust to missing observations. *Comm. Statist.-Simulation Comput.* **15**, 345-363.
- Andrews, D. F. and Herzberg, A. M. (1979). The robustness and optimality of response surface designs. *J. Statist. Plann. Inference* **3**, 249-257.
- Atkinson, A. C. and Cox, D. R. (1974). Planning experiments for discriminating between models (with discussion). *J. Roy. Statist. Soc. Ser. B* **36**, 321-348.
- Atkinson, A. C. and Donev, A. N. (1989). The construction of exact D-optimum experimental designs with applications to blocking response surface designs. *Biometrika* **76**, 515-526.
- Atkinson, A. C. and Donev, A. N. (1993). *Optimum Experimental Designs*. Clarendon Press, Oxford.
- Bolt, Beranek, and Newman Software Products (1992). D-optimal designs. *Statistical Appendices*. 3-1-3-29. Bolt, Beranek and Newman, Cambridge, Massachusetts.
- Box, G. E. P. and Draper, N. R. (1987). *Empirical Model-Building and Response Surfaces*. Wiley, New York.
- Box, G. E. P. and Hunter, J. S. (1961). The 2^{k-p} fractional factorial designs. *Technometrics* **3**, 311-351.
- Box, G. E. P., Hunter, W. G. and Hunter, J. S. (1978). *Statistics for Experimenters*. Wiley, New York.
- Cook, R. D. and Nachtsheim, C. J. (1980). A comparison of algorithms for constructing exact D-optimum designs. *Technometrics* **22**, 315-324.
- Cook, R. D. and Nachtsheim, C. J. (1982). Model robust, linear-optimal designs. *Technometrics* **24**, 49-54.
- Cook, R. D. and Nachtsheim, C. J. (1989). Computer-aided blocking of factorial and response-surface designs. *Technometrics* **31**, 339-346.
- Fedorov, V. V. (1972). *Theory of Optimal Experiments*. Translated and edited by W. J. Studden and E. M. Klimko. Academic Press, New York.
- Franklin, M. F. (1985). Selecting defining contrasts and confounded effects in p^{n-m} factorial experiments. *Technometrics* **27**, 165-172.
- Galil, Z. and Kiefer, J. (1980). Time- and space-saving computer methods, related to Mitchell's DETMAX, for finding D-optimum designs. *Technometrics* **22**, 301-313.
- Harville, D. A. (1974). Nearly optimal allocation of experimental units using observed covariate values. *Technometrics* **16**, 589-599.
- Heavlin, W. D. and Finnegan, G. P. (1993). Adapting experiments with sequentially processed factors. *ASA Proceedings, Section of Physical and Engineering Sciences*, 44-49.

- Herzberg, A. M. and Andrews, D. F. (1976). Some considerations in the optimal design of experiments in non-optimal situations. *J. Roy. Statist. Soc. Ser. B* **38**, 284-289, with correction, **39**, 394.
- John, P. W. M. (1971). *Statistical Design and Analysis of Experiments*. Macmillan, New York.
- Johnson, M. E. and Nachtshiem, C. J. (1983). Some guidelines for constructing exact D-optimum designs on convex design spaces. *Technometrics* **25**, 271-277.
- Jones, B. (1976). An algorithm for deriving optimal block designs. *Technometrics* **18**, 451-458.
- Kacker, R. N. and Tsui, K.-L. (1990). Interaction graphs: graphical aids for planning experiments. *J. Quality Technology* **22**, 1-14.
- Läuter, E. (1976). Optimal multipurpose designs for regression models. *Math. Oper. Statist.* **7**, 51-68.
- Li, W. W. and Wu, C. F. J. (1997). Columnwise-pairwise algorithms with applications to the construction of supersaturated designs. *Technometrics* **39**, 171-179.
- Li, W. W. (1995). An algorithmic approach for the construction of optimal balanced designs. Technical Report, Carlson School of Management, University of Minnesota.
- McLean, R. A. and Anderson, V. L. (1984). *Applied Factorial and Fractional Factorial Designs*. Marcell Dekker, New York.
- Mitchell, T. J. (1974a). An algorithm for the construction of "D-optimal" experimental designs. *Technometrics* **16**, 203-210.
- Mitchell, T. J. (1974b). Computer construction of "D-optimal" first-order designs. *Technometrics* **16**, 211-220.
- Mitchell, T. J. and Miller, F. L. (1970). Use of design repair to construct designs for special linear models. Report ORNL-4661. Oak Ridge National Laboratories, Oak Ridge, Tennessee.
- Nachtshiem, C. J. (1989). On the design of experiments in the presence of fixed covariates. *J. Statist. Plann. Inference* **22**, 203-212.
- Nguyen, N.-K. (1994). Construction of optimal block designs by computer. *Technometrics* **36**, 300-307.
- Park, J. S. (1994). Optimal Latin-hypercube designs for computer experiments. *J. Statist. Plann. Inference* **39**, 95-111.
- Radson, D. and Herrin G. D. (1995). Augmenting a factorial experiment when one factor is an uncontrollable random variable: a case study. *Technometrics* **37**, 70-81.
- Riordan, J. (1958). *An Introduction to Combinatorial Analysis*. Wiley, New York.
- Sacks, J., Welch, W. J., Mitchell, T. J. and Wynn, H. P. (1989). Design and analysis of computer experiments. *Statist. Science* **4**, 409-435.
- Silvey, S. D. (1980). *Optimal Design*. Chapman and Hall, New York.
- Sun, D. X. and Wu, C. F. J. (1994). Interaction graphs for three-level fractional factorial designs. *J. Quality Technology* **26**, 297-307.
- Taguchi, G. (1987). *System for Experimental Design: Engineering Methods to Optimize Quality and Minimize Cost*. UNIPUB/Kraus International, White Plains, New York.
- Van Schalkwyk, D. J. (1971). On the design of mixture experiments. Ph.D. Thesis, University of London.
- Wu, C. F. J. and Chen, Y. (1992). A graph-aided method for planning two-level experiments when certain interactions are important. *Technometrics* **34**, 162-175.
- Wynn, H. P. (1970). The sequential generation of D-optimum experimental designs. *Ann. Math. Statist.* **41**, 1655-1664.

Advanced Micro Devices, MS 117, P. O. Box 3453, Sunnyvale, CA 94088-3453, U.S.A.

E-mail: bill.heavlin@amd.com

E-mail: paul.finnegan@amd.com

(Received September 1995; accepted September 1997)