Design of Experiment Algorithms for Assembled Products

CHRISTINE J. SEXTON

Office of National Statistics, Titchfield, Fareham, England, PO15 5RR

DAVID K. ANTHONY

Instituto de Acústica (CSIC), C/ Serrano 144, 28006 Madrid, Spain

SUSAN M. LEWIS, COLIN P. PLEASE, and ANDY J. KEANE

University of Southampton, Southampton, England, SO17 1BJ

Designing experiments to identify improvement in products that are assembled from manufactured components does not readily fit into conventional design of experiments methods and can be costly. Efficient methods are explored for determining designs for engineering problems where some, or all, of the factors of interest are (a) not easily set to prescribed values and (b) are dependent on a combination of properties of several components. The methods involve taking a sample of each type of component, measuring the relevant features and then finding a design that specifies an optimal set of assembled products for experiment. Three examples from manufacturing industry are presented to illustrate the approach. Two different algorithms for finding designs are described, an exchange algorithm and a genetic algorithm, and a comparison of their performances is made on the three examples.

Key Words: *D*-Optimality; Exchange Algorithm; Genetic Algorithm; Search.

Quality improvement experiments in industry commonly make use of fractional factorial or response surface designs to investigate the effects of factors on product performance. These designs cannot be used when the nature of the product, or its manufacturing process, imposes practical restrictions on the combinations of the factor levels that can be run. In such experiments, computer search algorithms are often employed to find a suitable design and these algorithms are widely available as commercial software. This paper addresses the difficulties that arise in the design of experiments for manufactured products and, in particular, for mechanical products that are assembled from several components. In general, some, or all, of the factors to be investigated have values that depend on a number of features of the components. Many of the factors may not be readily set to particular experimental values so that conventional design methods are not applicable. This paper presents and compares algorithms that allow designs to be found efficiently for experiments on such products. The practical problem is that sets of components are available that have had their relevant features measured and it is now required to select components and combine them into assembled products in such a way that maximum information can be obtained on the factors of interest.

Dr Sexton is Assistant Methodologist in the Office for National Statistics. Her email address is c.j.sexton@ons.gsi.gov.uk.

Dr. Anthony is Research Fellow at the Instituto de Acústica (CSIC), Madrid. His email address is iaca344@ia.cetef.csic.es.

Dr. Lewis is Professor of Statistics in Southampton Statistical Sciences Research Institute. Her email address is S.M.Lewis@soton.ac.uk.

Dr. Please is Professor of Applied Mathematics in the School of Mathematics. His email address is cpp@maths.soton.ac.uk.

Dr. Keane is Professor of Computational Engineering in the School of Engineering Sciences. His email address is ajk@soton.ac.uk.

Two different algorithms are presented, one an exchange algorithm and the other a genetic algorithm, and their performances are compared on three practical examples from engineering companies.

For experiments where there are constraints on the design region, such as mixture experiments (Cornell (1990)), exchange algorithms (EAs) are the most popular method of searching for optimal designs. These algorithms search iteratively for improved designs by replacing points in a current design with points selected from a list of permitted or candidate points; see, for example, Myers and Montgomery (2002), pp. 413-414, for a description of EAs. In these algorithms, the most commonly used design selection criterion is D-optimality. This criterion is appropriate when accurate estimation is required of the coefficients in the linear model that is assumed to describe the response when the experiment is being planned; typically a low-order polynomial model. In the following work, we also assume a low-order polynomial for the response with independent, normally distributed errors having constant variance.

Genetic algorithms (GAs) may also be used to search for efficient designs. These algorithms work on a large pool, or generation, of designs that are represented by chromosomes; that is, strings of numbers that encode the values of the factors. The fitness, or performance, of each design in the generation is evaluated, for example, using D-optimality. Two stages are used to develop designs for the next generation. First, in the selection stage, each of the good designs is chosen with high probability but each poor design may also be chosen with a low probability. Second, from the selected designs, the next generation of designs is created by crossover and mutation procedures. In a crossover procedure, features of pairs of designs are combined, whereas, in mutation, random changes are made to a design. Crucial to the effectiveness of a GA is the choice of values for the parameters that control each procedure and the coding or representation of the factor values in the chromosomes. GA methods have been found to compare favourably with EA methods by Heredia-Langner et al. (2003), particularly for finding designs when linear restrictions exist on the design region for the experiment. An alternative approach of simulated annealing (see Haines (1987)) has also been used successfully for similar problems.

Whether the design selected for an experiment is a standard factorial design, a response surface design or a design obtained from a search algorithm, it will specify combinations of factor values that must be achieved with reasonable accuracy in order for the experiment to produce useful results. When manufactured products are assembled from several components, the production of products that meet the specifications of the chosen design may be prohibitively expensive due to, for example, the cost of manpower to measure large samples of components or the cost of fabricating special components.

One approach in such circumstances is to take large samples of components and measure their relevant properties. Subsamples of these components are then identified that provide a set of assembled products that best explore the design region of interest for a required experiment size. This basic approach is due originally to Harville (1974), who considered a number of different design problems, including an experiment from the chemical industry on batches of dyestuff, where the value of a covariate, impurity of a raw material, was known for each batch. More recently, O'Neill et al. (2000) applied a similar technique to the investigation of chemical treatments for preventing wood cracking after weathering, where the uncontrollable properties of the wood samples were measured and included as noise factors. The methods presented in this paper enable designs to be found for this type of experiment as well as for more complicated design problems on assembled products.

Similar issues arise in experiments to identify the tolerance levels to be set in the manufacture of components. One approach taken is to have a few special components made and to use these more than once by reusing them in different products or runs in the experiment; see Shainin (1993) and Bisgaard (1997) for detailed methods. An alternative approach, presented by Bisgaard et al. (2000), applies to products for which performance can be simulated using a CAD system so that experiments may be performed computationally. The methods given in this paper are applicable where reuse of components is impossible and no simulation package is available.

A further complication in designing experiments for assembled products is that the factors that are thought likely to influence product performance may not be directly measurable. Rather, the factors may have values obtained as functions of several parameters from one or more components. Such factors are called *derived factors* by Sexton et al. (2001). An example is the journal clearance of a bearing, which is derived as the difference between the outer diameter of a shaft and the inner diameter of a hole in a bearing block into which it fits. Unless all the parameters defining the factor can be fully controlled and set, conventional design of experiments cannot be applied.

The following method of experimentation is considered for these practical problems. Samples of components of the various types are obtained from production and, after careful measurement of the features of interest, a design is found that specifies the arrangement of these components into a set of assembled products for the experiment. The combinations of factor values that correspond to each of the assemblies of components may be obtained directly from the measurements on the component features, or, in the case of derived factors, by calculation of the appropriate functions of the measurements. Careful choice of the selection and arrangement of components for assembly is required in order to obtain as much information as possible on the factors under investigation. The possible combinations of factor values that can be used in the experiment depend on the component measurements available from the samples. Thus, there cannot be full control over the factor values used in the experiment. In practice, experiments may include parameters that can be set to required values, as used in conventional factors, as well as parameters that can only be measured, as illustrated below. We have assumed that the levels of such conventional preset parameters can be chosen independently of the measured parameters, as this is the only situation we have encountered to date.

The purpose of this paper is to illustrate industrial experiments on assembled products and to use three examples to compare two algorithms for searching for efficient designs. The recommended algorithm has been incorporated into a software package called DEAP (Design of Experiments for Assembled Products), which is freely available for download from www.maths.soton.ac.uk/~sml/screen_assemble, that enables users to investigate particular assembled products and manufacturing processes. The software is written in Visual C with a user interface for Windows machines.

Industrial Examples

The methods and algorithms described below were developed in response to experimental design problems arising in industry. The following gives an outline of three problems to which the methods have been applied and the particular challenges that each problem presented. Detailed descriptions of these problems, issues arising in practice, and designs and statistical analyses have appeared elsewhere in the literature, as indicated below. In addition, the problems feature as worked examples in the step-by-step software guide, DEAP (2005), and the data sets to enable others to work through examples and compare methods and approaches are available from the website.

Example 1: Hydraulic Gear Pump (Pilot Experiment)

The first example is an initial study on a hydraulic gear pump, manufactured by Sauer-Danfoss Inc., which is assembled from seven components: a flange, two bearing blocks, a drive gear, a driven gear, a housing and a cover (see Figure 1). Hydraulic fluid is drawn through the pump by rotating gear teeth. The efficiency of the pump is reduced by fluid leaking around the internal components, which is caused by the high pressure in the fluid. An experiment was required to determine those aspects of the pump that have an important influence on the leakage, where leakage is measured by the difference between the theoretical and actual flow rates through the pump. Of primary interest in the initial study were three leakage paths that were expected to be important. These paths were identified with the three physical quantities: the journal clearance, the gear form (a function of the profile of the meshing gear teeth), and the side gap (a measure of the distance between the gears and the bearing blocks). The values of these three leakage paths could not be measured directly but they could be derived from measured geometrical features of the components, such as the widths of the gears, the bearing blocks and the housing.

A conventional factorial experiment could not be used to investigate these derived factors because, for example, the gear teeth could not be machined to the required accuracy within acceptable costs. The alternative approach of the detailed measurement of a large sample of gear teeth in order to select a specified number with the required dimensions was prohibitively expensive. Further, the reuse of components within the experiment was not a viable option because testing the pump involves the gear teeth cutting into the housing and bearing blocks to produce a tight seal, and this produces permanent changes in relevant features of the components. Thus, this example considered three derived factors, namely, the three leakage paths, with all other aspects of the pumps held as fixed as possible. For the pilot

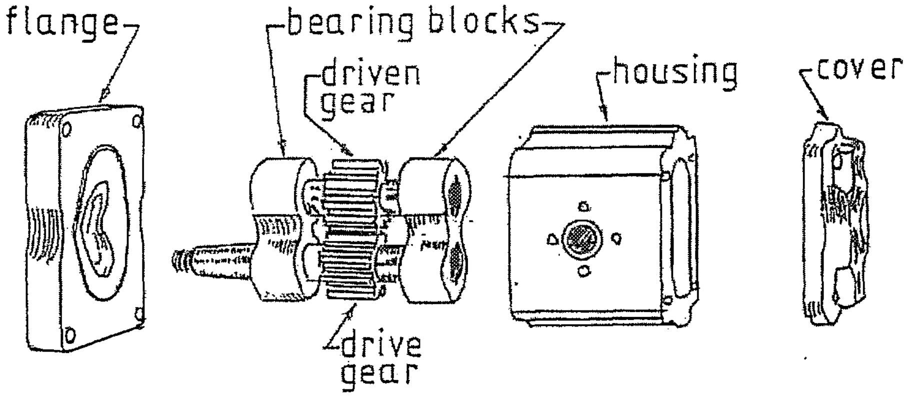

FIGURE 1. Hydraulic Gear Pump.

study, 12 driven gears, 24 bearing blocks and 11 drive gears were available and their features were measured. This resource allowed the assembly and testing of 11 pumps.

The full details of this experiment, including the component measurements, values of the derived factors, the final plan used and the leakage measurements are available in Sexton et al. (2000) and Sexton et al. (2001), with the former paper concentrating on engineering practicalities and the latter emphasizing statistical analysis methods. Further statistical issues related to power and the size of experiments, and illustrated using data from this study, are discussed by Sexton and Lewis (2001).

Example 2: Hydraulic Gear Pump (Follow-Up Experiment)

The second example concerns a further experiment on the pump of Example 1, which built on the findings from the pilot study and examined some additional factors. The results from the initial study indicated that the gear form should be investigated further. This was achieved by considering two separate elements of this derived factor that represented different aspects of the leakage path. Thus, gear form was replaced by two derived factors (involute form and lead-edge error) in the follow-up experiment. The derived factor side gap was retained for further investigation and a new derived factor, the fit of the bearing in the housing, was also included. There was no evidence from the pilot study that the journal clearance factor was important and hence this factor was not included in the follow-up experiment.

From a physical understanding of the working of the pump, it was decided to include three additional aspects of the pump design in this second experiment. The factors were the horizontal position of the flange and of the cover and the endfloat of the gear pack in the housing. All these factors could easily be set to prescribed levels. The available engineering knowledge indicated that the effect of each factor on the leakage was likely to be quadratic, with a minimum expected to occur within the range of interest. Hence, three levels of each of these factors were used in the experiment.

This example thus investigated four derived factors and three conventional factors, which could be set to prespecified levels. Components sufficient to make 44 hydraulic gear pumps were available for the experiment.

The plan used, together with the values of the conventional and derived factors, and the leakage responses were given by Sexton et al. (2000) and Sexton et al. (2001). The details of the measured components, the derived factors and the process used to investigate and generate the design for this example form the basis of the Advanced Tutorial in DEAP (2005). The tutorial also covers practical aspects, such as defining products, components and lists of parts, together with setting up the database of measurements associated with these, and the functional dependence of the factor values on these measurements. In addition, the tutorial indicates how to obtain simple graphical summaries of the measurements that allow, for example, anomalous values to be identified.

Example 3: Electro-Acoustic Transducer

An experiment was planned to investigate the sound output from an electro-acoustic transducer

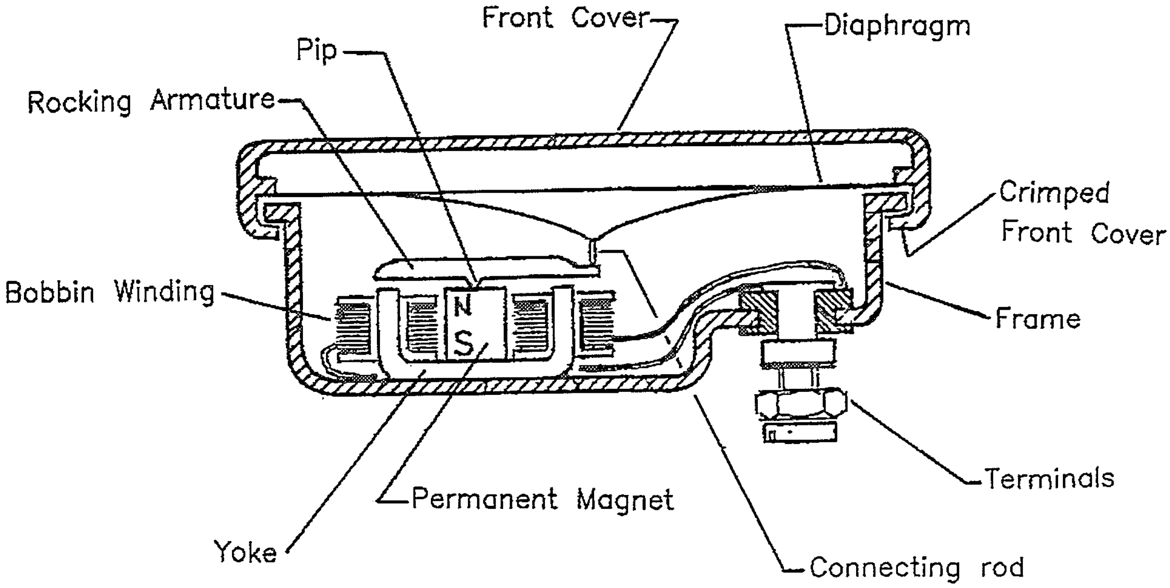

FIGURE 2. Schematic of an Electro-Acoustic Transducer.

manufactured by Hosiden Besson Ltd. and illustrated in Figure 2. In this device, an oscillating input electrical signal is fed to the bobbin windings, which causes the armature to rock on the pivot created by the pip. One end of the armature is connected to a lightweight diaphragm, which then radiates sound. The air gaps between the rocking armature and the yoke need to be carefully adjusted so that the balanced state of the armature is maintained during operation. This adjustment is made manually by an operator, who is guided by an electrical measurement (the electrical-gap adjustment). The magnetic field strength of the permanent magnet is then adjusted by another operator (in a process known colloquially as demagging) to achieve the required sound output level from the capsule for a given driving electrical signal.

Two of the factors to be explored, the magnetic permeability of the yoke and of the armature, could be measured but could not be produced to specified values. A total of 139 yokes and 113 armatures were available, together with their measured magnetic permeabilities. However, these yokes and armatures experienced variations in the heat-treatment process (through batch variation and kiln position) and hence heat treatment factors were also included in the design of the experiment.

There were four conventional factors to be studied in the experiment: the asymmetry of the bobbing windings, the electrical-gap adjustment, the skill of the demagging operator and the crimping force applied to hold the front cover to the capsule frame. In addition, there were four factors that were associated with measured components, namely, the magnetic permeabilities of the armature and of the yoke, the heat treatment batch and the armature pip height. The resources allocated to this experiment allowed a total of 60 capsules to be assembled.

Details of the factors and their levels, together with the plan and an analysis of the results are given by Anthony et al. (2003). This example is included in the user documentation, DEAP (2005), where full details are available, including the measurements.

Design Algorithms

Two competing types of algorithm were developed to search for designs among the large number of possible sets of assemblies of components that could be used in an experiment. The first is based on an exchange algorithm (EA) and the second is a genetic algorithm (GA).

Within each algorithm, a criterion is required for choosing between designs. In the examples in this paper, the main aim of each experiment was to understand how the factors jointly influenced the measured response. Hence, the D-optimality criterion was chosen which minimizes the generalized variance of the estimated model coefficients. The objective function to be minimized is $D = (|\mathbf{X}'\mathbf{X}|/n^p)^{-1/p}$, where n is the number of runs in the experiment, p is the number of coefficients in the assumed model and $\mathbf{X}'\mathbf{X}$ is the information matrix of the design (see, for example, Myers and Montgomery (2002), p. 394).

The Exchange Algorithm

The algorithm presented here differs from conventional exchange algorithms for design search in two ways. It is able to find designs when factor values, and their combinations, can only be selected from those available from the samples of components with the restriction that each component can only be used once. In addition, the algorithm allows designs to include factors, the values of which depend on several measured properties of different components.

This section provides an outline of the algorithm. First, the procedure for obtaining a starting design is described. Second, the steps taken to improve on this, and subsequent, designs are summarized.

Suppose that each component has been uniquely labeled and the relevant component variables measured. The exchange algorithm begins from a starting design that is obtained by a random allocation of parts from the set of measured components and from the possible preset factor levels to make the required number, n, of products for testing. For each of the assemblies in the starting design, the corresponding value of each factor is either (i) the preset level or (ii) the component measurement or, for derived factors, (iii) an engineering-based function of several measurements and preset levels. The n combinations of factor values form the starting design from which the first value of D is calculated.

There are three stages of improvement on the starting design. In broad terms, stage I attempts to improve the design by making changes to the combinations of the levels of the preset factors while holding fixed the combinations of factor values that involve component measurements. Stages II and III make changes to the combinations of factor values that arise from component measurements while keeping the preset factor combinations fixed. The objective function D is calculated at every stage from the combinations of values of all the factors in the design.

- Stage I: Search conventional factors. The Modified Fedorov exchange procedure of Cook and Nachtsheim (1980) is used to find the optimal set of preset factor combinations. The allocation of the measured components is held fixed at this stage.
- Stage II: Search unallocated measured components. The conventional preset factor values are held fixed at this stage and at stage III.

  1. Starting with the first product, as defined by the combination of components in the design, each component is swapped with each component of the same type that is not yet allocated to a product. The swap that leads to the greatest improvement in D is accepted.
  2. Step 1 is repeated for each type of component.
  3. Steps 1 and 2 are repeated until the improvement in D is smaller than some prespecified value.

- Stage III: Search by reallocation of measured components between products.

  1. Starting with a particular type of component, each component of this type within the design is interchanged with every other component of the same type within the design. The interchange that leads to the greatest improvement is retained.
  2. Step 1 is repeated for each of the remaining types of component in the product.
  3. Steps 1 and 2 are repeated until the improvement in D is smaller than some prespecified value.

The steps of Stages I, II and III are iterated until the improvement in the D value is sufficiently small. The tendency for the search to become trapped in a local optimum is overcome by making multiple runs, or tries, of the algorithm from different starting designs in the usual way.

Stages II and III replace the procedure in the original algorithm given in Sexton et al. (2001), where improvement in the design was sought by selecting $k > 1$ combinations of components in the design, and systematically substituting into the design all possible new sets of k component combinations, selecting the set that gave the greatest improvement in the value of D. The replacement of this procedure by Stages II and III has substantially improved the performance of the algorithm in terms of its speed and the frequency with which good designs are found.

Not all of Stages I, II and III are relevant to all product assembly problems. Example 1 requires only Stages II and III, as there are (i) no conventional factors at Stage I, and (ii) one driven gear and two bearing blocks are available for swapping at Stage II. Example 2 requires only Stages I and III, as there are no surplus components to be considered, while for Example 3, all three stages need to be used.

In applications of the Modified Fedorov exchange algorithm, as used in Stage I, it is usual to allow iterations to continue until the change in D is very small, for example 10^{-4}. However, in the applications presented in this paper, the rate of convergence of the overall method was significantly improved by undertaking only one iteration of this algorithm each time Stage I was performed.

The Genetic Algorithm

Genetic algorithms were first described by Holland (1975) and have been popularised mainly through the work of Goldberg (1989). These methods have been applied to design of experiment problems by, for example, Heredia-Langner et al. (2003), who gave an accessible introduction to the basic ideas of GAs. For product assembly problems, it is necessary to modify the basic approach, to decide on suitable coding for the factors and to determine appropriate values for the parameters within the various selection, mutation and crossover procedures (known as tuning the GA). A brief overview is given here; full details of the various approaches investigated and comparisons of the resulting speeds achieved are given by Anthony and Keane (2004).

In order to take account of the restrictions on the factor combinations that can be used in the design, a permutation-based GA was used. The set of measured components was coded in the GA using chromosomes that consisted of concatenated permutations, where each permutation represented a possible allocation of the components of one particular type into the assembled products. If there are more components of one type than are needed to form the products to be tested, then only part of the corresponding permutation is required. Also, if each product requires more than one component of a particular type as, for example, for the bearing block in Example 1, then that permutation is subdivided so that each subdivision indicates the particular instance of an individual component within a product.

The main difference between the GAs developed here and the algorithms available previously in the literature is the method used to encode the actual component allocation into the chromosome and the resulting modifications to the GA procedures. Two different coding methods were considered. The first method, *step index permutation*, used a set of $N-1$ indices to represent a permutation of N components. It has the advantages that there is no redundancy in the chromosome representation and that the standard operations used for binary representations for site selection for the crossover operation can be used. The mutation operation is slightly more complicated, however, in order to ensure that invalid permutations do not result. The second coding method, *extended chromosome representation*, used a simple listing of each permutation but consequently required special operators for the crossover and mutation procedures; see Goldberg (1989). This second method had some redundancy in the chromosome, as only $N-1$ integers are required to define a permutation of N integers.

The efficient application of GAs, particularly to nonstandard problems, requires careful coding and identification of the many parameters that define the operation of the algorithm. As the parameter values may be problem specific, it is important to balance the computational effort expended in tuning the GA against the gain in the speed of convergence. For each of the examples considered here, the optimal GA parameter values and a choice of coding method was made by performing fractional factorial experiments involving both quantitative and qualitative factors having 2, 3 and 4 levels, in which the GA was only taken part way toward convergence. This approach is similar to, but simpler than, factorial experiments that have been used to find good GA parameters, for example, by Pongcharoen et al. (2002) and Lee and Fan (2002).

Application and Comparison of the Exchange and Genetic Algorithms

Before the performances of the two design search algorithms were compared on the three industrial problems, the GA and the EA were applied to two conventional design problems in which all the factors were fully controllable. This allowed the GA to be benchmarked against a Modified Fedorov procedure, as only Stage I of the exchange algorithm was needed. The first of these, Example A, had three 3-level factors and four 2-level factors, while the second, Example B, had seven 3-level factors. In each of these examples, a full second-order model was assumed for the response and the number of runs was 10 more than the number of model coefficients: Example A had 42 runs and Example B had 46 runs. For each algorithm, 10 tries were made and the design with the smallest value of D was chosen. The results showed that, for both example A and B, the

TABLE 1. Values of D after 100,000 Evaluations

Example	1	2	3	A	B
Genetic algorithm	4.55	6.97	3.50	1.479	2.051
Exchange algorithm	4.50	5.21	3.16	1.490	2.118
GA/EA relative efficiency (%)	99	75	90	100.7	103.2

GA converged more quickly than the EA, and that it also produced designs with smaller D values. However, the D values were not substantially improved; see Table 1, with further details provided by Anthony and Keane(2004).

The GA and the EA were each used to find designs for the three industrial problems outlined earlier in this paper. Each algorithm was tried 10 times from a different starting design of a random allocation of components to products and, where appropriate, a random choice of combinations of the preset factor levels. Each time D was evaluated, as the algorithms proceeded, the smallest value of D found so far was recorded. The average of the smallest D values found from the 10 tries was plotted against the number of D evaluations; see Figures 3–5. This approach is similar to plotting the average best determinant against the objective function evaluations of Heredia-Langner et al. (2003). As mentioned in that paper, a comparison of the number of evaluations is not equivalent to a comparison of the time taken to run the algorithms because the EA can make use of updating procedures for calculating determinants that are not available to the GA.

Figure 3 shows the performance of the two algorithms when applied to Example 1. For this example, a design was required for three derived factors, with seven model coefficients in a second-order model and 11 runs in the experiment. For this small design, both algorithms converge in a small number of evaluations. There is little difference between the two methods, although the EA appears to converge slightly faster. After about 10,000 evaluations of D, little further improvement is evident.

The experiment of Example 2 has four derived fac-

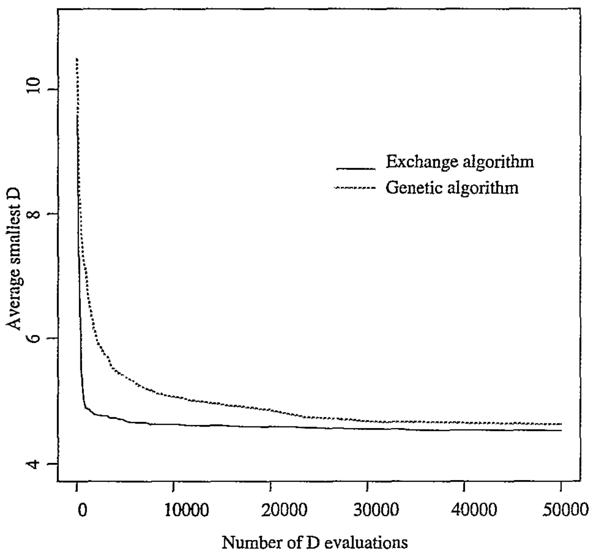

FIGURE 3. The Average, over 10 Tries of Each Algorithm, of the Smallest D Value Found after a Given Number of Evaluations of D for Example 1.

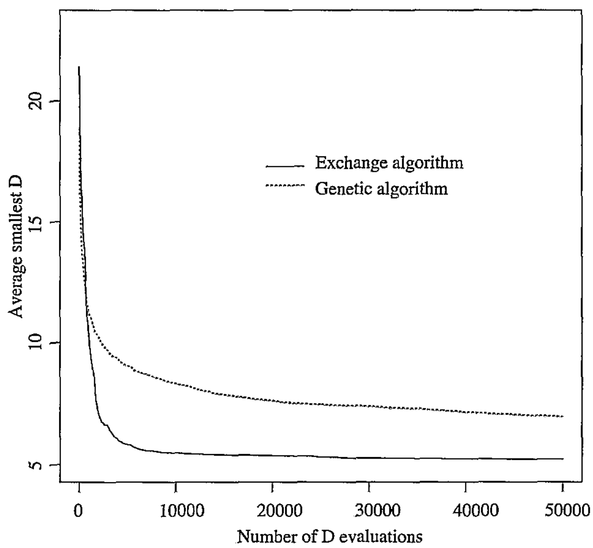

FIGURE 4. The Average, over 10 Tries of each Algorithm, of the Smallest D Value Found after a Given Number of Evaluations of D for Example 2.

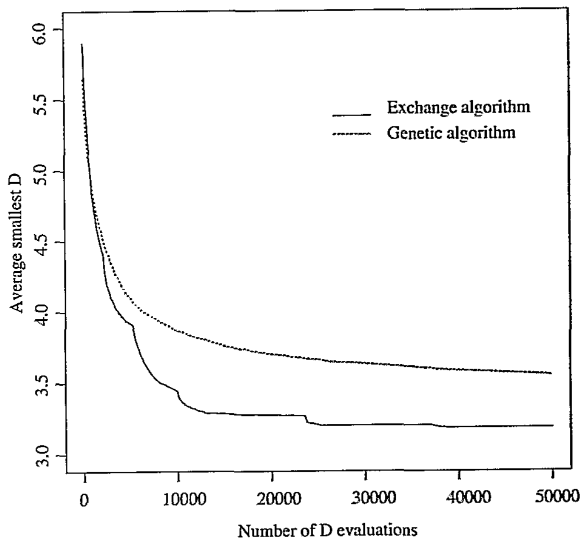

FIGURE 5. The Average, over 10 Tries of each Algorithm, of the Smallest D Value Found after a Given Number of Evaluations of D for Example 3.

tors and three conventional factors. A second-order model with 36 coefficients was assumed for the response and a design for 44 runs was required. Figure 4 shows that, initially, the GA converges with slightly fewer evaluations but that, ultimately, the EA achieves small values of D in fewer evaluations.

Example 3 has four conventional factors and four factors the values of which depend on the particular set of components, with a second-order model with 36 coefficients assumed to describe the response. A design was required with 60 runs. Figure 5 shows that, again, the GA gives an initial rapid drop in the D value but, subsequently, converges more slowly.

Table 1 summarizes the final value of D achieved after each of the algorithms had been used for a total of 100,000 D evaluations on each example. This shows that, for all three experiments, the EA finds a design with a lower value of D. For Examples 1 and 3, the difference in the final D value is not substantial. The difference is greater for Example 2, where the GA is 25% less efficient than the EA.

Conclusions and Discussion

Two different search algorithms have been applied to find designs for three industrial experiments to improve the engineering design and manufacturing process of a product. In each case, the product has properties that preclude the use of standard methods to design the experiments. The methods presented address the problems associated with experiments on assembled products where some factors cannot readily be set to specified values, samples of measured components are used in the experiment, with each component used only once, and some factor values depend on measurements from more than one component.

For these examples, the EA performed well relative to the GA, although in the early stages of the search, the GA converges more quickly than the EA for complex problems. This behaviour was consistent, although the three examples varied considerably in the number of factors and factor levels. The two methods were also tested on two conventional problems, and here the GA performed better than the EA, with the difference increasing as the complexity of the design problem increased, in agreement with the findings of Heredia-Langner et al. (2003).

Possible reasons for the different relative behaviour of the algorithms for the two types of problems are the following. As discussed by Heredia-Langner et al. (2003), at any iteration in the algorithm, the EA search is constrained to consider only points in a candidate list, whereas a GA has no such limitation. In searches for conventional designs with quantitative factors, there is no intrinsic reason why the points in the design should be restricted and hence the GA is able to slightly out perform the EA. This can be viewed as the global searching of the GA being more efficient than the local searching of the EA.

For the assembly problems discussed in this paper, the search is limited to those sets of combinations of factor levels that are obtainable from the available components. Thus, although initially the global searching of the GA is more successful, the constrained local searching of the EA gives an eventual advantage.

It is possible that, for some very large problems, convergence to an optimal D value may be speeded up by combining the two algorithms so that the GA is used at the early stages before switching to the EA to complete the search. This approach has not been investigated. However, an efficient switch-over point from GA to EA could be identified by investigations similar to those used to optimize the parameter settings for the GA.

In practice, a design for an experiment would not be chosen on the basis of a single criterion such as

D-optimality. Other properties of possible designs should be examined, such as variance inflation factors and the variance–covariance matrix of the model coefficient estimators. Such issues were considered when applying the final designs to the industrial examples by using a short list of designs found from the algorithms, but have not been presented here.

The EA algorithm was implemented within DEAP to guide the design of experiments for assembled products. This decision was made because the performance of the EA and GA algorithms was similar across the various problems; however, the EA did not require the complication of automating an initial investigation in order to tune the parameters for each particular problem.

The algorithm in DEAP has been further extended to allow experiments to be run as an integral part of the manufacturing process through the use of sequential experimentation. At any time in the experiment, a pool of measured components is available. As each new set of components is added to the pool, the algorithm determines which components should be removed for assembly into the next product to be tested prior to shipping. In our current experience, economic considerations require the pool to be small and so an exhaustive search can be used to determine the optimal assembly to supplement those products built in the earlier stages of the experiment. Insight into the performance of the product is thereby gained over a period of time. This approach is particularly useful when products are produced at a low rate or few can be made available for testing at any one time.

Acknowledgments

This work was supported by the Engineering and Physical Sciences Research Council grant GR/N16754, together with Jaguar Cars, Hosiden Besson and Goodrich Engine Control Systems. We are grateful to the referees and the editor for helpful comments.

References

ANTHONY D. K. and KEANE, A. J. (2004). "Genetic Algorithms for Design of Experiments on Assembled Products". Technical Report No. 385, School of Mathematics, University of Southampton, Southampton, SO17 1BJ, UK.

ANTHONY, D. K.; SEXTON, C. J.; CHARLSTON, W. B.; KEANE, A. J.; LEWIS, S. M.; and PLEASE, C. P. (2003). "An Efficient Experiment Methodology to Investigate Product Design: An Acoustic Sounder Case Study". *Proceedings of the International Conference on Managing Innovative Manufacturing*, 22–23 September 2003, Aalborg, Denmark, pp. 291–302.

BISGAARD, S. (1997). "Designing Experiments for Tolerancing Assembled Products". *Technometrics* 39, pp. 142–152.

BISGAARD, S.; GRAVES, S.; and SHIN, G. (2000). "Tolerancing Mechanical Assemblies with CAD and DOE". *Journal of Quality Technology* 32, pp. 231–240.

COOK, R. D. and NACHTSHEIM, C. J. (1980). "A Comparison of Algorithms for Constructing Exact D-Optimal Designs". *Technometrics* 22, pp. 315–324.

CORNELL, J. A. (1990). *Experiments with Mixtures: Designs, Models, and the Analysis of Mixture Data.* Wiley, New York, NY.

DEAP USER MANUAL (2005). University of Southampton, Southampton, UK, SO17 1BJ. Available at http://www.maths.soton.ac.uk/~sml/screen_assemble

DUPPLAW, D.; BRUNSON, D.; VINE, A. E.; PLEASE, C. P.; LEWIS, S. M.; DEAN, A. M.; KEANE, A. J.; and TINDALL, M. J. (2004). "A Web-Based Knowledge Elicitation System (GISEL) for Planning and Assessing Group Screening Experiments for Product Development". *ASME Transactions—Journal of Computing and Information Science in Engineering* 4, pp. 218–225.

GOLDBERG, D. E. (1989). *Genetic Algorithms in Search, Optimisation and Machine Learning.* Addison-Wesley, Cambridge, MA.

HAINES, L. M. (1987). "The Application of the Annealing Algorithm to the Construction of Exact Optimal Designs for Linear Regression Models". *Technometrics* 29, pp. 439–447.

HARVILLE, D. A. (1974). "Nearly Optimal Allocation of Experimental Units Using Observed Covariate Values". *Technometrics* 16, pp. 589–599.

HEREDIA-LANGNER, A.; CARLYLE, W. M.; MONTGOMERY, D. C.; BORROR, C. M.; and RUNGER G. C. (2003). "Genetic Algorithms for the Construction of D-Optimal Designs". *Journal of Quality Technology* 35, pp. 28–46.

HOLLAND, J. H. (1975). *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control and Artificial Intelligence.* The University of Michigan Press, Ann Arbor, MI.

LEE, L. H. and FAN Y. L. (2002). "An Adaptive Real-Coded Genetic Algorithm". *Applied Artificial Intelligence* 16, pp. 457–486.

MYERS, R. H. and MONTGOMERY, D. C. (2002). *Response Surface Methodology: Process and Product Optimization Using Designed Experiments.* Wiley, New York, NY.

O'NEILL, J. C.; BORROR, C. M.; EASTMAN, P. Y.; FRADKIN, D. G.; JAMES, M. P.; MARKS, A. P.; and MONTGOMERY, D. C. (2000). "Optimal Assignment of Samples to Treatments for Robust Design". *Quality and Reliability Engineering International* 16, pp. 417–421.

PONGCHAROEN, P.; HICKS, C.; BRAIDEN P. M.; and STEWARDSON, D. J. (2002). "Determining Optimum Genetic Algorithm Parameters for Scheduling the Manufacturing and Assembly of Complex Products". *International Journal of Production Economics* 78, pp. 311–322.

SEXTON, C. J.; DUNSMORE, W.; LEWIS, S. M.; PLEASE, C. P.; and PITTS, G. (2000). "Semi-Controlled Experiment Plans for Improved Mechanical Engineering Designs". *Proceedings of the Institute of Mechanical Engineers, Part B* 214, pp. 95–105.

SEXTON, C. J. and LEWIS, S. M. (2001). "Determination of the

Size of Semi-Controlled Experiments with Derived Factors". *Statistics and Applications* 2, pp. 25–39.

SEXTON, C. J.; LEWIS, S. M.; and PLEASE, C. P. (2001). "Experiments for Derived Factors with Application to Hydraulic Gear Pumps". *Journal of the Royal Statistical Society, Series C* 50, pp. 155–170.

SHAININ, R. D. (1993). "Strategies for Technical Problem Solving". *Quality Engineering* 5, pp. 433–448.

COPYRIGHT INFORMATION

TITLE: Design of Experiment Algorithms for Assembled Products
SOURCE: J Qual Technol 38 no4 O 2006
WN: 0627402184005

The magazine publisher is the copyright holder of this article and it is reproduced with permission. Further reproduction of this article in violation of the copyright is prohibited.

Copyright 1982-2006 The H.W. Wilson Company. All rights reserved.