Proceedings of DETC'03
ASME 2003 Design Engineering Technical Conferences and
Computers and Information in Engineering Conference
Chicago, Illinois, USA, September 2-6, 2003

DETC2003/CIE-48275

WEB-BASED KNOWLEDGE ELICITATION AND APPLICATION TO PLANNED EXPERIMENTS FOR PRODUCT DEVELOPMENT

David Dupplaw
Faculty of Mathematical Studies
University of Southampton

David Brunson
Jaguar Cars Limited

Anna-Jane E. Vine
Faculty of Mathematical Studies
University of Southampton

Colin P. Please
Faculty of Mathematical Studies
University of Southampton

Susan M. Lewis
Faculty of Mathematical Studies
University of Southampton
Southampton SO17 1BJ
ENGLAND
Email: S.M.Lewis@maths.soton.ac.uk
Tel: +44 (0)2380 593652

Angela M. Dean
Department of Statistics
Ohio State University

Andy J. Keane
School of Engineering Sciences
University of Southampton

Marcus J. Tindall
Health Care Research Unit
University of Southampton

ABSTRACT

When planning experiments to examine how product performance depends on the design, manufacture and environment of use, there are invariably too few resources to enable a complete investigation of all possible variables (factors). Here we describe a web-based system for eliciting company knowledge into an efficient two-stage group screening method. The method investigates the effect of a large number of factors by grouping them in a first stage experiment whose results identify factors to be further investigated in a second stage. Central to the success of the procedure is ensuring that the factors considered, and their grouping, are based on the best available knowledge of the product. We present a web-based software system that allows information and ideas to be contributed by engineers at different sites and uses these expert opinions to guide decisions on the planning of group screening experiments. The software includes elements that predict the total resource needed for the experiment It also simulates the results of the experiment and estimates the likely percentage of important or active factors that fail to be detected. The approach is illustrated through the planning of an experiment on engine cold start optimization at Jaguar Cars.

INTRODUCTION

In multi-national companies where expertise is distributed across many centres, it is important to be able to access and pool existing knowledge and experience in order to guide the effective and efficient design of new and existing products and manufacturing processes. Design improvements can be identified by use of planned experiments (design of experiments or DOE) to determine the crucial factors in the design. A first step in such problems is therefore to identify a list of possible factors that might influence product performance. Often, in practice, a very large number of contending factors is identified for possible investigation and this makes direct application of conventional DOE methods impractical. The existing knowledge pool within the company concerning these factors needs to be accessed and interrogated in order to optimise the competing risks of using an overly large experiment which wastes resources against using too small an experiment which may produce confusing results. The aims are to eliminate from consideration those factors whose influence is believed to be negligible and to guide decisions on how to investigate the remaining factors most efficiently. In this paper we present a web-based software system that implements a methodology to identify factors and extract knowledge about their action by use of a dynamic questionnaire which can be used iteratively. The system allows users to summarise and explore knowledge to inform and direct the planning of efficient screening experimentation. The paper illustrates the use of this system to plan experiments for engine cold start optimisation at Jaguar Cars.

be investigated or taken into account in the planning of experiments.

Disadvantages of such brainstorming sessions include the fact that the dynamics of the meeting, such as the personalities or positions of authority of the participants, can affect the identified lists of factors. A further source of possible unintentional bias is that the convener of the meeting has to identify the participants in advance. An important role of the methodology that is presented here is to circumvent some of these difficulties and also to provide a focused starting point and further stimulus for informed debate in such meetings.

Questionnaires are widely used to gather information, for example in surveys, and to inform the use of Bayesian methods [1, 2] of data analysis. An alternative approach to eliciting expert opinions is given by Grimshaw, et al [3]. Our approach uses a questionnaire to collect knowledge about factors and exploits web-based technology to allow the questionnaire to evolve as new ideas are contributed. The information gathered is used to plan the experiment and can also be useful when the data from the experiments are to be analysed using Bayesian methods [4].

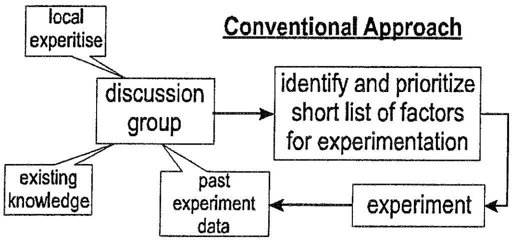

Figure 1. Conventional approach to gathering information

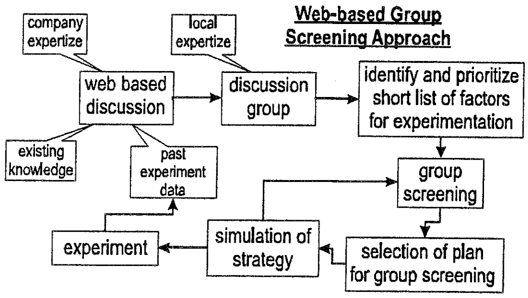

Figure 2. Elicited expert data-led approach to information gathering

BACKGROUND

The conventional approach to gathering information to inform the planning of experimentation is through the use of "brainstorming" techniques (see figure 1). In this approach a set of people with relevant expertise is identified and brought together for a focused meeting. The purpose of the meeting is for ideas to be put forward and a consensus reached on the factors for investigation and the role that the factors might play in the performance of the product. The end result of such a meeting would be (i) a list of factors that should be varied during the experiment, because they are believed to be important for product performance, and (ii) a list of factors that should be kept constant to ensure that any performance variability that they might cause does not mask the impact of the other factors. In addition, any available knowledge and opinions would be pooled on joint action (interaction) between, typically, pairs of factors that need to

Our approach has been developed as a result of our experiences with more conventional information gathering procedures for planned experiments. Where design improvement is primarily in the domain of a small group of people, we have observed that important factors may be overlooked. This can be due to the group of people bringing very similar experience and knowledge to the problem, and can be overcome by combining the views of diverse groups of people, most importantly from both the design and the manufacturing teams (see figure 2). We have found that the use of a conventional questionnaire overcomes some of the problems involved in eliciting opinions from a wide group of people. However, the fixed structure of the questionnaire exerted too much control on the format of the answers elicited and, by its nature, excluded the possibility of generating dialogue. We have tried out conventional questionnaires and evaluated their use with collaborating industries in the aeronautical, agricultural, electromechanical and automotive fields where the planned experiments took many different forms.

In this paper we apply a web-based approach to inform the planning of a type of screening experiment known as two-stage group screening, discussed by Lewis and Dean [5], and describe a specific application in the automotive industry. The aim is to find those factors whose main effects and interactions (Meyers and Montgomery [6], page 87) are sufficiently large to produce a substantive improvement (in engineering terms) in the mean and variation of the product performance. The minimum size of this worthwhile improvement is judged by engineers on the basis of scientific and cost returns on improvement to the performance.

In screening experiments, factors are typically investigated at two values or levels, called the "high" level (believed likely to produce the highest performance or response across the factor range to be explored) and a "low" level (where the lowest response is anticipated). In group screening, these individual factors are first put in groups and the factors within each group are varied *together* between these two levels. This enables new "grouped factors" to be defined, each of which represents a group of factors.

This type of screening involves two stages of experiment. At the first stage, an observation is made on each of the various different combinations of the levels of the grouped factors according to an experiment plan, for example a fractional factorial plan [6]. From an analysis of the data (such as estimation of main effects and interactions or a Bayesian analysis), a decision is made on which of the groups are important. The individual factors within these groups are then taken forward to a second stage experiment whose aim is to identify the key individual factors that affect product performance.

Two different two-stage group screening strategies are available in the software system described here; for technical descriptions see [5,7]. In one strategy, only main effects are investigated at the first stage (*classical group screening*) whilst, in the second, interactions between pairs of factors as well as main effects are examined at the first stage (*interaction group screening*). In product development experiments, the factors are usually of two types: *control or design factors*, whose values can be set during the manufacturing process, and *noise factors* whose values vary as they arise from the environment where the product is used, or from variation in the manufacturing process. Noise factors can often be set, or mimicked in an experiment. The technique takes account of the fact that some interactions, particularly those between control and noise factors, are of greater interest than others in improving product performance. In forming groups of factors, only factors of one type (control or noise) are put in a group in order to allow examination of interactions between grouped control and noise factors.

A key question when an experiment is being planned is whether the resources it will require are economically feasible. In two-stage group screening, the total experiment size cannot be determined with certainty when the experiment is planned due to the two-stage nature of the study. Information about the likely importance of the factors, elicited from experienced engineers, can be used to derive the probability distribution of the predicted total number, S, of individual main effects and interactions to be estimated. This number serves as a surrogate for the total experiment size.

In order to decide on the strategy, the number of groups of factors and how many factors to put in each group, the following practical criteria may be used:

(a) minimize the expected total experiment size
(b) minimize the probability of exceeding a budgeted target for experiment size
(c) maximize the probability of detecting the important design factors, especially those which can be set to make product performance less sensitive to varying noise factors.

In the software system described here, criteria (a) and (b) are implemented using the distribution of S, and criterion (c) by simulation of two stage group screening experiments.

METHODOLOGY

The process of setting up a plan for an experiment on a number of physical products is often a recursive process. Hence a system to plan such experiments needs to be flexible and interactive. In addition, the steps in the process need to be documented so that the information that was used to support any new understanding can be properly examined subsequently. The strategy employed is to use as much information as possible to guide the grouping of factors in a two-stage group-factor experiment. The procedure is not meant to be algorithmic but rather to act as a tool to inform decisions on possible grouping strategies in light of the particular product being considered. Through an analysis of the processes underlying the planning of experiments, the following five phases have been identified:

1. **Setup** - The experiment organiser determines the aspect(s) of the product to be investigated in the experiment. In particular, the primary measure of "performance" of the product, together with those factors within the product that may affect the performance. Each factor is identified as a control or noise factor and, for each factor, a range of values is suggested. Finally, three additional pieces of information are required to guide the subsequent decisions on planning the experiment. These are i) a target for the total number of individual observations on the experimental products that can be used in the experiment based on available time and resources, ii) an estimate of the minimum change in performance measure that would be considered to be of practical value, iii) an estimate of the likely error in the measured

observations, available from previous experiments or pilot studies.

2. **Acquisition of Information** - In this phase the above information supplied by the organiser is extended through the collection of views from other experts in the field. There is a need to obtain assessments of the importance of each factor in affecting the product performance, and views on the direction in which such influence may act as the factor level changes. Both quantitative and qualitative information may assist at this point, as well as an indication of each responder's confidence in their views. It is also important that any additional factors that have been overlooked, but which might influence product performance, are identified. When such new factors are identified, assessments of the directions of their effects and any likely interactions need to be made. This information gathering activity is conventionally achieved by a brainstorming session. In our software system, we have expanded this idea by using a questionnaire which allows more people to be involved and for each of them to give carefully considered comments before they are discussed in the more active forum of a brainstorming session. It also allows those completing the questions to add new factors that may then, in turn, be commented on by others in an iterative fashion.
3. **Summary of Importance** - The information gathered from the experts, from the questionnaire responses and from the brainstorming sessions, needs to be summarised to inform decisions concerning the planned experiment. In particular, the importance anticipated for each factor needs to be examined so that a probability can be assigned to its main effect being active. (This is needed for the assessment of possible groupings and strategies using criteria (a) and (b) above). Typically some factors will be expected to have a minimal effect on the performance, some to have a substantial effect and there will be others where there is much less certainty about their influence on performance. Typically, the organiser may decide not to investigate factors which are judged by the experts to be of little importance in order to reduce the experiment size, but he/she may include those factors where the experts are uncertain about the factor's importance. This stage of the process must draw on both the quantitative information, in the form of spreadsheet data, and also more qualitative information in order to draw up the final list of control and noise factors.
4. **Choice of Groupings and Strategy** - When the factors to be investigated in the experiment have been identified, the various different strategies and group sizes for the two stage experiment need to be assessed and compared. For a particular strategy (classical or interaction group screening) and particular choices for the number of groups and the numbers of factors put into each group, the probabilities assigned to the individual factor main effects are used to assess the total expected experimental effort that would need to be expended to examine i) the effects of the grouped factors as well as ii) the effects of the individual factors within those groups of factors found to have important effects. The approach is completely general in that it allows any groupings of the factors to be investigated under each of the group screening strategies and for any choice of probabilities that the individual main effect and interactions are active. By considering a number of different groupings, ranging from a small number of large groups to a large number of small groups, a short list of economical choices can be formed using criteria (a) and (b) above for each of the strategies. It may be necessary to iterate the procedure at this point, perhaps by reassessing which factors to include in the experiment, if it is found that an experiment cannot be carried out within the available resources.
5. **Simulation** - Having identified a short list of groupings and strategies which can be expected to fit within the resources on using criteria (a) and (b), the user can make an informed judgement on the basis of the proportions of active effects that are likely to be missed (criterion (c)) using simulation of the two stage experiments. The simulation enables two important practical concerns to be addressed: (i) where the direction of effects is unknown or incorrectly assessed, it is possible that a grouping may inadvertently be made in such a way that the contributions to product performance from the factors within a group cancel each other out, and (ii) within a group, several very small contributions from factors may combine to mask the effect of an active factor having a different direction. The simulation software calculates the probability of failing to detect important factors due to such possible cancellations.
   To simulate the outcome of using each of the possible groupings would be highly computationally intensive. Hence simulation is performed on the short list of grouping strategies drawn up by considering the predicted experiment size. It is then necessary to decide if one of the strategies is appropriate, that is, likely to lie within available resources and likely to have acceptably small proportions of the active main effects and interactions that are undetected. If a suitable strategy has not been found, then an increase in the resources for the experiment or a reduction in the number of factors investigated may have to be considered.

The above five steps guide the planning of a two stage group screening experiment from gathering available expert knowledge, to assessing a short list of plans against available resources for the experiment and an acceptable level of risk of failing to detect important effects. Several iterations of the steps are often needed in order to reach a final decision on the plan.

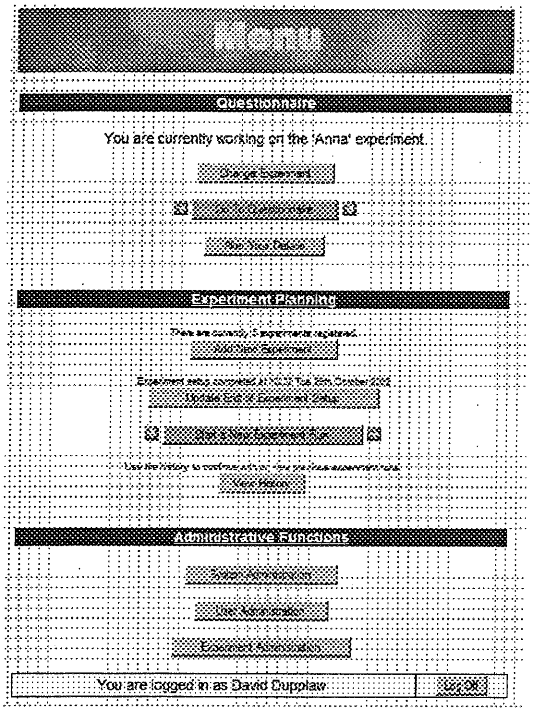

Figure 3. Front page of Group Screening software

SOFTWARE IMPLEMENTATION

The five steps described above have been incorporated into a software system which allows the necessary information to be collected and organised in an interactive manner. The design of the flow of the software system reflects the flow of this experimentation process with the need to return to various steps at any time being an integral part of the procedure.

The system resides on a central server where all the information is stored and calculations performed. Currently the system has been tested on both Linux and Windows XP operating systems. The system has been carefully constructed to provide transparent access to users of the system. All interaction with the system is through a web browser and user browsers tested include Internet Explorer, Netscape Communicator and Opera. The system has a number of different levels of user, namely administrators - who can initiate and analyse investigations into new product problems, users - who can input information and analyse various strategies, and guests who can only input information to particular problems.

The system is based on open source software available over the internet. This provides a cost effective way to provide much of the functionality required.

Installation of the software has to be carried out by a network administrator or support staff. It requires the installation of four pieces of software:

Apache web server - to provide the web based interactive environment that allows the users to be physically dispersed.
MySQL database server - to provide the database structures needed to control the system and to collect and store the information within the system.
PHP Tools and code - to provide the functionality and flexibility of the system.
Group screening and simulation code - to perform the necessary calculations needed in the assessment procedure.

All these tools are provided with the software but require root or administrator privileges to install.

Central to the software (whose menu is shown is figure 3) is the interactive questionnaire that elicits from users and guests their opinions on the importance of factors and their possible influence on product performance. Traditionally this would be done by a paper-based method, but this software allows the information to be gathered via the world-wide web, avoiding duplication errors, the labour of aggregating results, and printing and administration costs of the paper questionnaires.

The user is presented with a split screen which, on the left, displays a list of the possible factors in an experiment and, on the right, a panel where questions relating to each factor will be shown. When the user clicks on the name of a factor, a questionnaire is presented to the user in the right-hand panel. The questionnaire tries to quantify the user's understanding of the influence of the factor on the performance by asking the questions shown in figures 4 and 5.

The possible responses to the questions shown in the figures contain only approximate ideas of the response. If the user indicates a belief that varying the value of this factor will alter the performance of the product, the questionnaire asks for an indication of the expected trend of this response. Using "A" as the low value for the factor and "B" as the high value the user indicates whether the performance of the product is expected to greatly increase, or decrease, or only slightly increase, or decrease as the

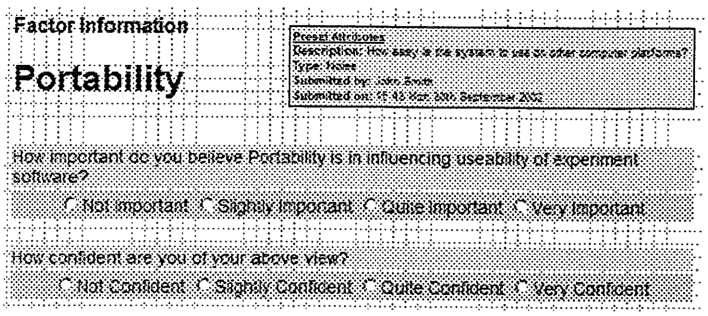

Figure 4. Page to indicate in the questionnaire the expected importance of a factor

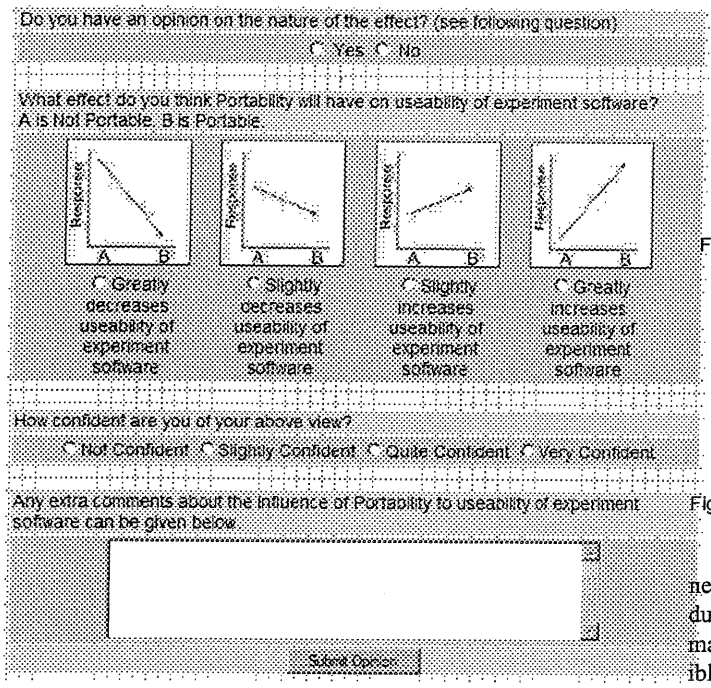

Figure 5. Page to indicate in the questionnaire the response trend of a factor

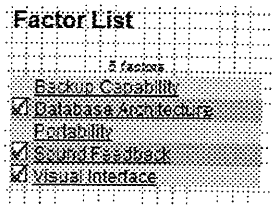

Figure 6. Page to indicate factors considered during the questionnaire

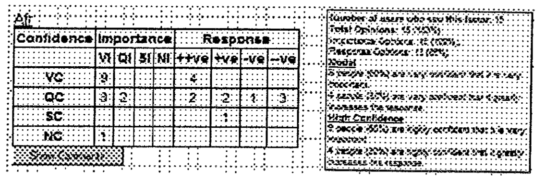

Figure 7. Page to summarise quantitative questionnaire data on a factor

factor changes from "A" to "B". In the current implementation, such trends are only shown as linear but, if a user considers the response to be non-linear, then the comment section allows them a written explanation of this opinion. For each question, the user selects from the radio buttons the answer which matches their opinion. The questionnaire responses are stored in a database, and are recalled when a user logs back into the system. This means that users can complete as much of the questionnaire as they initially have time for, and then come back at another convenient time to complete more. They may also alter their opinions by selecting from the factor list a previously completed opinion.

If a user addressing the questionnaire believes that a factor which is not listed may be important to the performance of the product, they may add the factor themselves by using the "Add My Own" button (not shown) at the bottom of the factor list. As they do this, a screen is presented into which both the factor name and a brief description of the factor and the possible levels to consider can be entered. On submission of this new factor, the factor list is updated and the questionnaire can be immediately completed for this new factor. The additional factor will be visible to any new, or returning, users.

As the user submits opinions with the "Submit Opinion" button, the factor list is updated on the left-hand side of the page. If the user correctly filled in the fields, then a small tick is placed next to the completed factor (see figure 6). If a mistake was made during the completion of the questionnaire (for example, not all mandatory fields were completed), then a small cross will be visible next to that factor. Factors where the current user has given no opinion will have no visible marking next to the factor name.

The summary of the data collected through the questionnaires is presented in several forms. At any time, users can view all opinions including written comments. However, to aid understanding of the quantitative data from the questionnaire, the data are summarised in small tables. These show not only the number of opinions given for each possible response direction, but also the confidence that the users expressed in their opinions. An example of such a table is given in figure 7. In the example, most responders were confident that the factor in question was very important. This display helps to give the administrator a clear view of the importance of each factor and hence allows an allocation of a probability to the factor being important for the later grouping stages. Since there is seldom sufficient information to allow any precise probabilities to be defined, the factors are categorised into one of the three categories i) Definitely affect response ii) May affect response iii) Negligible effect on response. Probabilities for these categories are then requested. Included in these probabilities are the probabilities that indicate the importance anticipated of the interactions between pairs of factors. These can either be entered manually or can be calculated using the method of Chipman [2].

The software then allows the user to assess both classical group screening and interaction group screening over two stages of experiment. Our experience of these methods indicates that, although interaction group screening usually requires more experimentation than classical, the risk of failing to detect active interactions can be far less, particularly for the important interactions between control and noise factors.

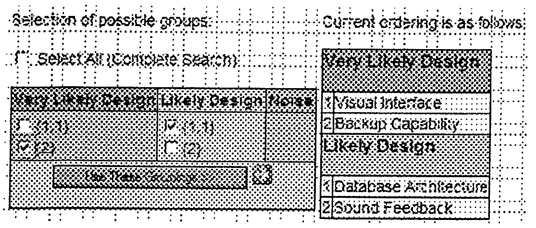

Figure 8. Page to indicate group sizes for assessment by software

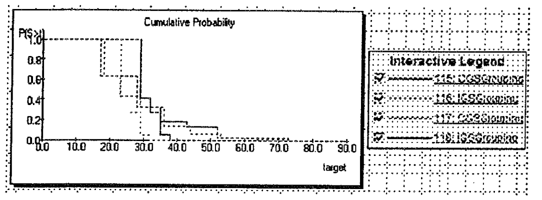

Figure 9. Page to compare four strategies for the probability of the total experiment size exceeding a given target size

For a given grouping of the factors, specified by the user as in figure 8, the software allows the user to calculate the probability of exceeding the available experimental resource for both strategies for any specified set of groupings. In order to simulate the two stage experiments the software requires plans for the first stage of experiment on the grouped factors. Currently, efficient fractional factorial plans which are tailored to prioritize the investigation of main effects and interactions between control and noise factors are provided by the software by default, with the option of a user-supplied plan being available. The results of the simulations are stored and, at any time, a user can review all the assessments of any particular grouping strategy. (The calculations can be fairly computationally intensive, so it is preferable to store results rather than re-run them). There is no automatic ranking of these strategies and the user is given several ways of viewing the data to aid in making their decision. Figure 9 shows an example of a comparison of four different strategies using criterion (b) using superimposed graphs of the probability of exceeding a target experiment size. Once some choice of number and sizes of groups is identified as good, the simulation software is then run repeatedly to calculate the proportion of experiments in which important effects are missed.

CASE STUDY

The software described above has been used to plan a relatively large experiment at Jaguar Cars. The experiment concerned the optimisation of cold start performance using spark plug resistance measurement on a particular engine type. Design of Experiment methods are commonly used by Jaguar Cars for product improvement and the methodology described in this paper fitted in well with their conventional working practices. The main difference here was to exploit the web-based software to enhance the conventional "brainstorming" sessions to guide factor selection and also to exploit the group screening methodology to allow more factors to be considered than would be possible by conventional DOE methods.

1. **Setup** - After suitable discussions, the performance measure to be used to assess the quality of the product was identified as the percentage of a given number of engine cycles with a "low" spark plug resistance. The expected resource available for the total experiment was anticipated to be in the region of one hundred observations on different engine configurations. A two percent change in performance measure was perceived as the minimum worthwhile change and a standard deviation of the performance measurements of about 0.8% was obtained from a pilot study. An initial list of factors for consideration for the experiment was drawn up based upon previous, smaller scale investigations and local knowledge. Each of these factors was also identified as being either a control or a noise factor.
2. **Information gathering** - User names for access to the web-based software questionnaire were distributed to relevant Jaguar, Land Rover and Ford centres in the UK and USA. Two iterations on the questionnaire were used with a discussion occurring after receipt of the initial returns because a number of new factors had been proposed. Subsequently, responders were asked to comment again particularly on these new factors. Of special value was the general comment section of the questionnaire as this led to the need to clarify the definition and levels of certain of the factors.
3. **Summary of importance** - A total of fifteen people contributed to the questionnaire which gave a greater range of input than was conventionally obtained. The data were summarised into the tables discussed earlier. In addition the data helped to guide the "brainstorming" sessions by identifying certain factors as consistently judged important, and hence not requiring discussion, and other factors as consistently viewed as irrelevant. The face-to-face sessions could therefore concentrate on the more controversial factors identified through the questionnaire and on whether they should be included in the experiment. It was decided that, from the full list of factors, there were twelve control factors and two noise factors that needed to be investigated in the experiment. Of these, five control factors were believed to be very likely to be active with confident information on the expected direction of their effects. Each of these factors had a value of 0.7 assigned to the probability that the main effect is active. Our software investigations have shown that the assessments based on the expected size of the experiments (see criteria (a) and (b)) are insensitive to this precise value with a value of 1.0 giving similar results. There was much

less consistency of opinion on the remaining control factors but, overall, there was some support for the view that they were likely to be active. However, little consistent information was available on the direction of their influence. It was anticipated that perhaps only one would actually be active so each of them were assigned a probability of 0.14 (=1/7). Finally, it was anticipated that at least one of the two noise factors would be active so a probability of 0.5 was assigned to each. The interaction probabilities were assigned using the approach of Chipman [2].

Control factors	Noise factors
AFR	Injector tip variation
Spark Time	Humidity / Temp.
Calibration	
Engine off timing	
Idle Speed	
Plug Type	
Injection timing	
Spark Advance	
Transient fuel with calibration	
Plug gap	
Variable valve timing	
Injector spray angle/direction	

Table 1. Factors for experiment from the questionnaire and discussions

4. **Grouping** - Many possible strategies for grouping the factors are available and were considered. It was found that groups of similar size produce experiments that are expected to be smaller than those with very disparate group sizes. Two approaches to grouping the noise factors were considered: a single group, and two separate groups. As an illustration, figure 10 shows the results of the grouping software for the case where the "very likely to be active" control factors are in two groups of sizes 2 and 3, "likely to be active" control factors are in four groups of sizes 2, 2, 2, 1, and the noise factors are in a single group. The figure displays the distribution of the predicted size of the experiment and the resulting risk of exceeding various targets. It was found that the experiment would fit within the available resources and, as expected, that classical group screening (CGS) would need considerably less resource than interaction group screening (IGS). Two economic groupings, for each method, were taken forward to be assessed further by simulation.
5. **Simulation** - As expected from earlier results, the simulation showed that the groupings with CGS risked missing a far higher percentage of active interactions between control and noise factors than the groupings with IGS. As these interactions were of particular interest, the IGS method was adopted. In choosing between whether the noise factors should be put into one or two groups, the results showed that there was little difference in the expected total experiment size but the proportion of missed active interactions between control and noise factor was much greater when the noise factors were in a single group. Hence the final grouping adopted was to use the noise factors in two sepa-

Probability of 5 very likely control main effects being active = 1.0 Chipman values 0.01, 0.125, 0.125, 0.5						
5 very likely control factors	7 likely control factors	2 noise factors	$E(S)$	$\sigma(S)$	Prob. $S > 110$	Prob. $S > 120$
Grouping						
2,3	2,2,2,1	2	102	7.57	0.15	0.01

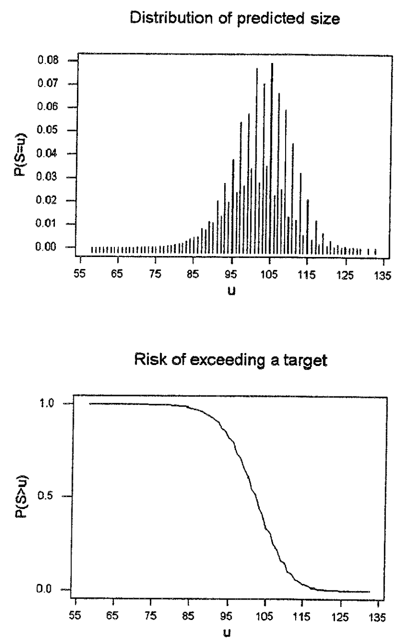

Figure 10. Information for assessment of a grouping for an experiment under interaction group screening.

rate groups. The fractional factorial plan used to simulate the experiment then provided the basis of the plan for use in the actual experiment. The final plan required randomization before it is used and this had to accommodate the practical constraint imposed by four of the factors whose levels were time-consuming to change.

CONCLUSIONS AND FUTURE WORK

The use of a web-based system was found to have considerable advantage over the conventional approach to information gathering to guide planned experiments. It enabled a international group of experts to pool their ideas and knowledge in order to plan an efficient and effective experiment to examine factors which are relevant to product improvement. A further advantage was that the software provided a record of the information underlying any new decisions made in changing a product design, as well as an archive of the new knowledge found from the experiments. The software also made accessible recent research developments in the area of group screening that allow investigation of a larger number of factors than can be accommodated through conventional approaches. Interaction group screening results in a larger experiment than conventional group screening. However, through investigation of a larger number of factors and interactions between them, the experiments enable more accurate identification of the important factors and a decrease in the probability of missing key interactions. Hence these methods are economically more efficient and reduce the chance of the results of the experiment being inconclusive. To ensure the software is adequately responsive for general use, improvements are being implemented in the speed of the algorithms. Also, extra functionality is being integrated to allow further elicitation of information, and greater user feedback. It is forseen that the software framework could be abstracted to allow greater flexibility for other application domains. Installing the system into any organization would be very straightforward and cost effective requiring only simple software installation. The two-stages of experiment outlined in this paper are scheduled to be performed at Jaguar Cars. Those wishing to try the system as beta-testers are welcome to contact Professor Susan Lewis for details.

ACKNOWLEDGMENT

This research was undertaken with support from The Engineering and Physical Science Research Council (Grant GR/N16754). Additional thanks are due to Jaguar Cars, Hosiden Besson, Goodrich Engine Control Systems for their support and collaboration on this work.

REFERENCES

[1] Box, G.E.P. and Meyer, R.D. "Finding the active factors in fractionated screening experiments", J. Quality Technology, 25, pp. 94-105, 1993.

[2] Chipman, H., "Bayesian variable selection with related predictors", The Canadian Journal of Statistics, 24, pp. 17-36, 1996.

[3] Grimshaw, S.D., Collings, B.J., Larsen, W.A. and Hurt, C.R. "Eliciting factor importance in a designed experiment", Technometrics, 43, pp. 133-146, 2001.

[4] Chipman, H., Hamada, M. and Wu, C.F.J, "A Bayesian variable-selection approach for analysing designed experiments with complex aliasing", Technometrics, 39, pp. 372-381, 1997.

[5] Lewis, S. M. and Dean, A.M, "Detection of interactions in experiments on large numbers of factors", J. Roy. Statist. Soc. B, 63, pp. 633-672, 2001, (with discussion).

[6] Meyers, R.H. and Montgomery, D.C. *Response Surface Methodology: Process and Product Optimisation using Designed Experiments*. Wiley, New York, 2002.

[7] Dean, A.M. and Lewis, S.M., "Comparison of group screening strategies for factorial experiments", Computational Statistics and Data Analysis, 39, pp. 287-297, 2001.