

READINGS IN
INFORMATION
VISUALIZATION
USING VISION
TO THINK

WRITTEN AND EDITED BY

STUART K. CARD
JOCK D. MACKINLAY
BEN SHNEIDERMAN

<i>Senior Editor</i>	Diane D. Cerra
<i>Director of Production and Manufacturing</i>	Yonie Overton
<i>Production Editor</i>	Cheri Palmer
<i>Production Assistant</i>	Pam Sullivan
<i>Assistant Editor</i>	Marilyn Uffner Alan
<i>Editorial Assistants</i>	Belinda Breyer and Sarah Luger
<i>Cover Design</i>	Ross Carron Design
<i>Cover Photo</i>	Michael McGovern/Nonstoc
<i>Text Design, Composition, and Pasteup</i>	Christine Cotting, Uppercase Publication Services, Ltd.
<i>Color Preparation</i>	Mark Ong, Side by Side Studios
<i>Copyeditor</i>	Gary Morris
<i>Proofreader</i>	Jennifer McClain
<i>Indexer</i>	Ty Koontz
<i>Printer</i>	Courier Corporation

Designations used by companies to distinguish their products are often claimed as trademarks or registered trademarks. In all instances where Morgan Kaufmann Publishers, Inc. is aware of a claim, the product names appear in initial capital or all capital letters. Readers, however, should contact the appropriate companies for more complete information regarding trademarks and registration.

ACADEMIC PRESS

An Imprint of Elsevier

525 B Street, Suite 1900, San Diego, CA 92101-4495, USA

<http://www.academicpress.com>

Academic Press

Harcourt Place, 32 Jamestown Road, London, NW1 7BY, United Kingdom

<http://www.academicpress.com>

Morgan Kaufmann Publishers

340 Pine Street, Sixth Floor, San Francisco, CA 94104-3205, USA

<http://www.mkp.com>

© 1999 by Academic Press

All rights reserved

Printed in the United States of America

07 06 05 5 4 3

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means—electronic, mechanical, photocopying, recording, or otherwise—without the prior written permission of the publisher.

Permissions may be sought directly from Elsevier's Science and Technology Rights Department in Oxford, UK. Phone: (44) 1865 843830, Fax: (44) 1865 853333, e-mail: permissions@elsevier.co.uk. You may also complete your request on-line via the Elsevier homepage: <http://www.elsevier.com> selecting "Customer Support" and then "Obtaining Permissions".

Library of Congress Cataloging-in-Publication Data

Card, Stuart K.

Readings in information visualization : using vision to think /
written and edited by Stuart K. Card, Jock D. Mackinlay, Ben
Shneiderman

p. cm.

Includes bibliographical references and index.

ISBN 1-55860-533-9

1. Information display systems. 2. Computer graphics.
3. Visualization. 4. Image processing. I. Mackinlay, Jock D.
II. Shneiderman, Ben. III. Title.

TK7882.16C36 1999

621.39'87—dc21

98-53660

CIP

Externalising Abstract Mathematical Models

Lisa Tweedie, Robert Spence, Huw Dawkes and Hua Su

Department of Electrical Engineering,
Imperial College of Science, Technology and Medicine
South Kensington, London, SW7 2BT
Tel: +44 171 594 6261
l.tweedie@ic.ac.uk

ABSTRACT

Abstract mathematical models play an important part in engineering design, economic decision making and other activities. Such models can be **externalised** in the form of Interactive Visualisation Artifacts (IVAs). These IVAs display the data generated by mathematical models in simple graphs which are interactively linked. Visual examination of these graphs enables users to acquire insight into the complex relations embodied in the model. In the engineering context this insight can be exploited to aid design. The paper describes two IVAs for engineering design: The Influence Explorer and The Prosection Matrix. Formative evaluation studies are briefly discussed.

KEYWORDS : Interactive Graphics, Visualization

INTRODUCTION

Many mathematical problems can benefit from being examined visually. Indeed most spreadsheets and statistical packages enable users to quickly create static representations of their data. These graphs have an accepted role as tools for mathematical problem solving. However the value of adding interactivity to such representations has yet to gain widespread recognition.

Responsive (i.e. rapid) interaction can facilitate active exploration of problems in a manner that is inconceivable with static displays. For example users can start to pose "What if" queries spontaneously as they work through a task. Such exploration can enormously facilitate the acquisition of qualitative insight into the nature of the task at hand, as well as revealing direct quantitative results.

In this paper we describe what we call Interactive Visualisation Artifacts (IVAs). These are environments developed to enable users to solve a particular task - in this case within the field of engineering design.

The IVAs we will discuss here differ from much existing work principally because we are not attempting to visualise

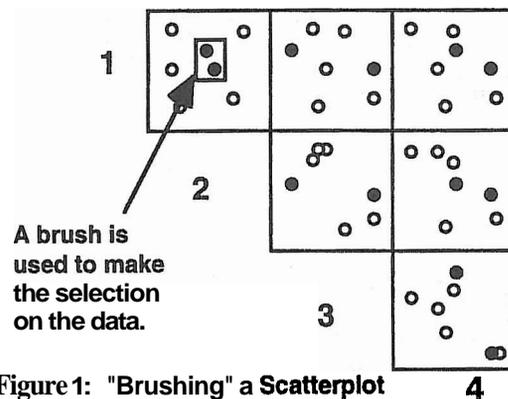


Figure 1: "Brushing" a Scatterplot

raw data but, rather, data which is precalculated or generated on demand from mathematical models. We also exclude data which maps comfortably onto natural representations e.g. 3D volumetric models of flow through a pipe. Instead we focus on more **abstract** mathematical models which have no obvious representation.

We can take as an example the design of an engineering artifact. Mathematical models (equations) exist which relate the artifact's performance to the parameters that describe the physical nature of that artifact. Thus, for a bridge, performances such as traffic capacity and cost can be calculated from a knowledge of parameters such as cable diameters and foundation depth. A designer needs to explore the relationships between parameters and performances in order to elicit a useful design.

The development of IVAs for such applications requires the creation of new representations that externalise pertinent aspects of the model. The IVAs we describe in this paper show how such novel representations can be created by **interactively linking simple graphs in several ways**. On a simple level we can link many similar graphs, as Becker et al [3] did with their "brushed" scatterplots (Figure 1). We can also link different **types** of representations together. For example, by selecting a subset of data on a histogram and colour encoding the same subset on a scatterplot. These links can also perform different functions - for example the selected subset could be colour encoded or it could be hidden from view.

Two IVAs for engineering design are described in this paper: the Influence Explorer and the Prosection Matrix. They exhibit powerful and effective linking both within and between IVAs.

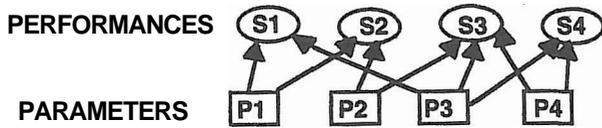


Figure 2: The Parameter->Performance relationship

Previous Work

The idea of linking graphical representations is not new. As early as '1978 Newton [12] was linking several scatterplots and colour encoding selections to discover trends in data. Many others have developed simple linking IVAs e.g. IVEE [1], Permutation Matrices [4], BEAD [5], SeeSoft™[6], AutoVisual [7], VisDB [10], Nested Histograms [13], The Table Lens [14], Visulab [15], The InfoCrystal [17], The Attribute Explorer [18] and The Dynamic HouseFinder [20].

Most of these IVAs only use one type of representation to display data. However a combination of representations may also be beneficial, since the user is then able to consider the problem from several different perspectives. Schmid and Hinterberger [15] have called this "Comparative Multivariate Visualisation" and embodied the concept in their "Visulab" software. Here four different representations (Parallel Coordinates [9], Andrews Plots [2], Permutation Matrices [4] and Multiple Scatterplots [3]) can be linked in several ways : encoding with colour, hiding part of the data and reordering the data. The use of several different representations of data, and the manner of their linking, is a key issue in the development of IVAs.

Visual Design Issues

The design of any IVA should proceed with various characteristics of visual problem solving in mind (Tweedie [19]). As Nardi and Zamer [11] point out, IVAs are external representations of the users problem which "stimulate and initiate cognitive activity". Zhang and Norman [21] identify that such external representations act as memory aids; provide information perceptually without need for interpretation; anchor and structure cognitive behaviour; and change the task.

Suchman [18] emphasises that "it is frequently only on acting in a present situation that its possibilities become clear". In other words users will often pick up information opportunistically from their environment. It is partly this tendency to stimulate opportunistic behaviour that makes IVAs interesting. Consequently, the visual cues provided must be designed to support this opportunistic process.

DESIGN FOR MANUFACTURABILITY

A typical task that has a mathematical model associated with it is that of engineering design. For a given product such as a light bulb, a model can be formed of the way the parameters (whose value is open to choice by the designer) influence performances (Figure 2). In the light bulb example, performances such as a bulb's brightness and its lifetime will partly be determined by parameters such as the number of coils in its filament and the thickness of that filament. The mathematical model is a set of equations, each relating a performance to a number of parameters.

The designer must choose numerical values for parameters in such a way that the performances they influence, usually in a very complex fashion, take on values acceptable to a customer. In other words, when designing a light bulb, the designer has to keep a specification in mind. If for example they are asked to design a light bulb that will be very bright and last for at least 6 months then they need to find the set of parameters values that will satisfy this specification.

The traditional design process

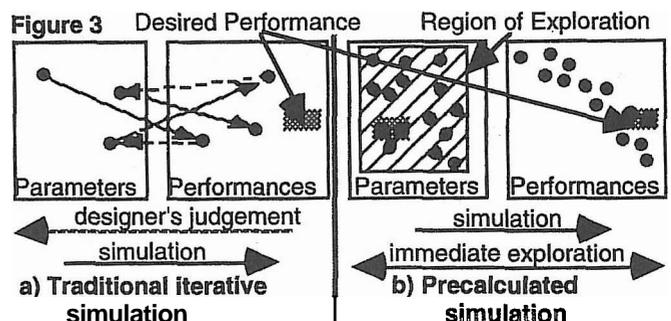
Given a set of parameters, an engineering artifact can be simulated to establish the corresponding performances. Unfortunately the reverse is not true: a designer cannot choose a performance value and calculate the parameters needed to achieve it. For this reason traditional design is characterised by a series of iterations in which the designer selects a set of parameters and then simulates the artifact to find out what the performances are. Design proceeds through the gradual adjustment of parameters until a satisfactory set of performance values is found. This design process is illustrated in figure 3a for an artifact defined by two parameters and influencing two performances. The design is represented by a single point moving in parameter and performance space. This "trial and error" approach can be tedious and time-consuming and is heavily dependent on a designer's expertise.

Precalculation

The design process can be immensely simplified if one has mathematical models of the relationship between parameters and performances. Figure 3b shows how such models can be used to create a precalculated exploration database. The designer selects a wide "Region of Exploration" in parameter space within which the final design might well be expected to lie. Within this region a large number of points (e.g. over 500) are generated randomly, each point representing a design. For each of these sets of parameter values the corresponding point in performance space is computed using the artifact's mathematical model: In our light bulb example, a dataset generated in this way would describe a variety of light bulbs each having randomly different parameter values and associated performances. The benefit of creating such a dataset is also illustrated in figure 3b. The designer can now readily select their desired performance values and "look up" which parameter sets give them those values.

Designing in the real world

Unfortunately the aim of engineering design is not simply that of finding a single set of parameter values that satisfies



a specification. Inevitable fluctuations in manufacturing processes mean that parameter values can only be guaranteed to lie within a so-called tolerance *range*. For example the filament width of our light bulb might vary slightly during manufacture, and this variation could have a crucial effect on a performance. We therefore need to define exactly how much each parameter can vary. The combined set of parameter tolerance ranges defines a *tolerance region* in parameter space. These are the bulbs that will be manufactured.

Figure 4 shows the rectangular tolerance region for the simple case of two parameters. In the same space, an irregularly shaped "Region of Acceptability" defines the location of all the artifacts that satisfy the performance requirements. Achieving a good design is a matter of fitting these two regions to each other with maximum overlap

Overall Design Objectives

As well as satisfying the customer's requirements on performance, it is usually the case that there is also some overall objective that must be achieved. One such objective is that of maximising the manufacturing yield, which is the percentage of mass-produced bulbs that satisfy the customer's requirements on performance. With reference to Figure 4, yield is that percentage of the tolerance region which lies within the region of acceptability.

Another such design objective might be the unit manufacturing cost of each bulb that is shipped to the customer. Usually the wider the tolerances are on the parameters the cheaper the bulb will be to manufacture.

THE INFLUENCE EXPLORER

Precalculation forms the backbone of the Influence Explorer. Once the data has been precalculated (as described earlier), it provides an exploration database on which to start an investigation. Figure 5 shows how the population of 600 precalculated designs is displayed in the form of histograms. All performance histograms are plotted horizontally to the left of the screen and the parameter histograms vertically to the right. An artifact is represented once on each plot in the appropriate bin. Each column in the histogram represents the number of designs that fall within that bin. In other words, the histograms are frequency plots.

Qualitative Exploration

In order to form an effective external representation of the task the Influence Explorer must allow the user to gradually

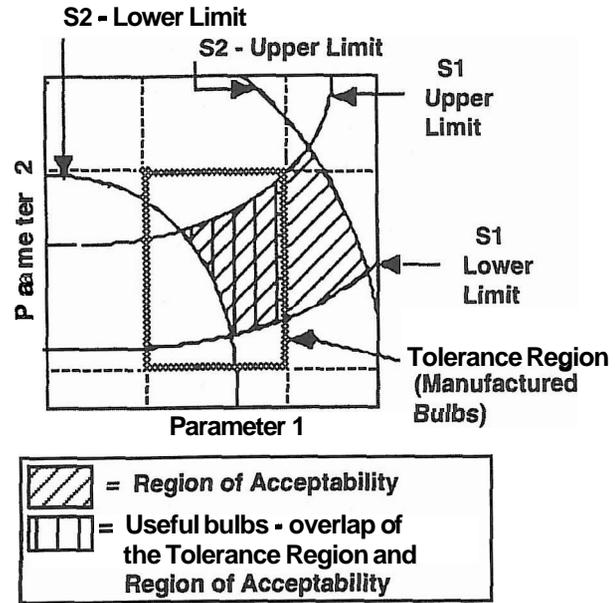


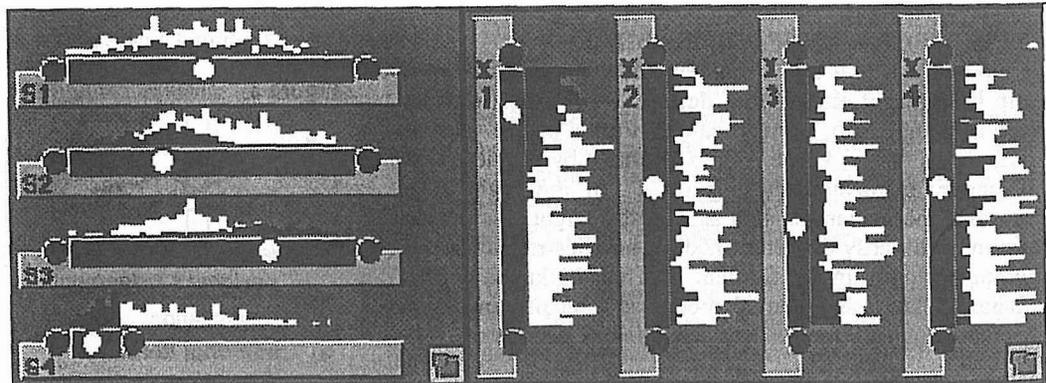
Figure 4: In 2-parameter space, two performances F1 and F2 define upper and lower limits for acceptable performance. Manufactured bulbs lie within the tolerance region.

build up a coherent picture of their problem, in other words the complexity must be introduced in stages.

In the initial stages of design the user will want to gain a *qualitative* understanding of the problem. The designer can place exploratory limits on parameters and performances, thereby defining ranges of those quantities. In Figure 5 a range of performance on S4 has been defined with a slider. This action leads to the colour linking (black) of those bulbs that lie within the selected range on the S4 histogram and all the other histograms, so that the selected subset can be viewed across all the histograms. The potential for exploring the inter-relation between parameters and performances is now apparent. Confidence in these perceived relations can be sought by interactively moving the selected range of S4 up and down its scale and observing the corresponding movement of the highlighted bulbs on the other scales. The power of such a dynamic action to generate insight is difficult to convey in static words and diagrams, but is *strikingly* obvious in actual use.

It is worth emphasising that the discovery of a "trade-off" relation between two performances is immensely important

Figure 5: The performance (left) and parameter (right) histograms. A selection has been made on S4 and these same points are highlighted on each of the other histograms. Circles indicate the mean of the selected points.



in engineering design. In the Influence Explorer this discovery is virtually immediate, whereas in conventional design practice such a trade-off might be discovered only after tedious search or, at worst, not at all.

Additional tools enhance the functionality of the Influence Explorer. A mouse-click on a bulb in one histogram highlights that same bulb, and displays corresponding values, in all the other histograms. Another option connects these points with a line and allows the comparison of several different bulbs. These lines are known as "parallel coordinate" plots [9]. Yet another option places a circle on each of the histogram scales indicating the mean of the currently selected bulbs (see figure 5). This is useful when a range is being moved as it eases detection of trends.

Quantitative Design decisions

As well as indulging in *qualitative* exploration, the designer must at some stage take note of the *quantitative* detail associated with a customer's requirements on performance. To do so a "specification option" is selected (Figure 6 - colour plate).

The placement of upper and lower limits on the performance scales invokes another linking mechanism. Red colour coding identifies bulbs that lie within all the performance limits, those bulbs which fail one limit are colour coded black, while dark and light grey denotes two and three failed limits respectively. Such colour coding provides valuable sensitivity information. For example, it is immediately noticed (Figure 6 - colour plate) that a relaxation of the upper limit on S4 would turn some black bulbs into (acceptable) red bulbs, knowledge which might well lead to a discussion about the wisdom of that particular upper limit. Negotiations concerning performance specifications are common to engineering and could be considerably clarified using this information.

Design for Manufacture

As already explained, inevitable variations in the manufacturing process are such that, in the design of a *mass-produced* artifact such as a light bulb, the designer must be concerned with the selection of parameter *ranges* rather than specific values. It is the combination of all these selected parameter ranges that must satisfy the performance limits defined by the customer.

Parameter ranges are defined by the selection of upper and lower limits (Figure 7 - colour plate), in exactly the same manner as for the performances. Again, the selection of parameter limits invokes a linking mechanism, once more leading to additional colour encoding. Though at first sight complex, the coding is, we suggest, matched to an engineering designer's real needs and, given the motivation provided by a tool offering responsive exploration, is readily, even eagerly learned. Figure 8 (colour plate) is a replica of Figure 4 with the relevant colour codings shown. Figure 8 and the table attached to figure 7 may help clarify the rationale behind this coding:

- Red denotes bulbs that satisfy all limits. They lie within parameter limits (and are therefore manufactured) and they satisfy the customer's performance limits.

- **Black** denotes a bulb that satisfies all the performance limits but lies outside one parameter limit, and is therefore not manufactured. Thus it will turn red if one parameter limit is adjusted to include it.

- **Blue** bulbs are those which are manufactured (and hence lie within parameter limits) but fail one or more performances. These are the bulbs which cause a reduction in yield. Tightening a parameter limit to eliminate blue bulbs (for example raising the lower limit of X1 in Figure 7) will reduce the number of manufactured artifacts which violate a customer's requirements, hence raising the yield. The Blue bulbs are coded in two shades of blue - **Dark Blue** indicates those bulbs that are manufactured and only violate one performance limit; relaxation of that performance limit will turn those bulbs into red ones (e.g. in figure 7 expanding the lower limit on S1 will turn the dark blue bulbs red). **Light Blue** indicates those bulbs which are manufactured and violate more than one performance limit.

- Grey bulbs are those which fail one parameter range and one or more performance limits. They would therefore turn blue if they were to be enclosed within the tolerance region. Thus in Figure 7 if the upper limit on X2 is extended to turn the black bulbs into red ones, this gain in the number of (red) acceptable bulbs would be offset by the number of grey bulbs turning blue and, thereby, adding unsatisfactory bulbs to the manufacturing process.

The principal **advantage** of such colour coding is that it indicates how *altering* the parameter or performance limits will effect the overall usefulness of the design.

Yield Enhancement

To facilitate design for maximum yield the Influence Explorer continuously computes, and displays in numerical form, the value of the yield. The designer may well begin by attempting to select parameter ranges that **maximise** the yield, hopefully to a value of 100%. In order to achieve such a high yield the user needs to adjust the tolerances taking account of where the red and therefore "useful" points lie and trying to reduce the number of blue points. By keeping an eye on the yield the user can slowly **optimise** their solution until they have found an optimum yield.

100% yield can obviously be achieved by **making** the parameter ranges sufficiently small (Figure 14 - colour plate), but another overall objective - the minimisation of manufacturing cost - militates against such a solution. It is normally the case that the wider the parameter ranges, the lower the cost of the artifact. There is therefore a strong incentive to select parameter ranges that are as wide as possible commensurate with an acceptably high yield (see Figure 15 - colour plate).

Focused Sampling

Unfortunately when interacting with tolerances limits the precalculated data set becomes a constraining factor in the Influence Explorer. Since the requirements are now becoming specific, it is unlikely that many of the original **600** points will fall within *all* the performance and parameter requirements. This curse of dimensionality results in very few colour coded points. To overcome this problem the Influence Explorer is programmed to

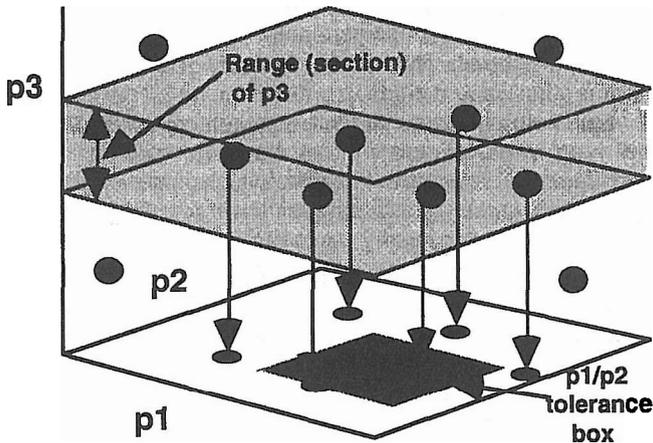


Figure 9: A section of p_3 is projected onto a p_1/p_2 scatterplot

dynamically resample the model so that a number of points always fall within and close to the tolerance region. Evidence of this process can be seen in Figure 7 where the column heights within the tolerance limits are higher than column heights on the rest of each parameter histogram.

THE PROSECTION MATRIX

The Prosection Matrix provides an alternative perspective of the model. It is a set of scatterplots (Figure 10) arranged in a matrix, as suggested by Becker et al [3]. Each scatterplot corresponds to a different pair of parameters, and all possible parameter pairs are represented. Thus, for the bulb's four parameters there are six scatterplots.

The construction of each scatterplot is illustrated conceptually in Figure 9 for the simple case of a 3-parameter system. p_1 and p_2 are the scatterplot's two parameters. p_3 is a third parameter on which a parameter range has been set. Only data that falls within p_3 's chosen parameter range is projected down onto the p_1/p_2 plane.

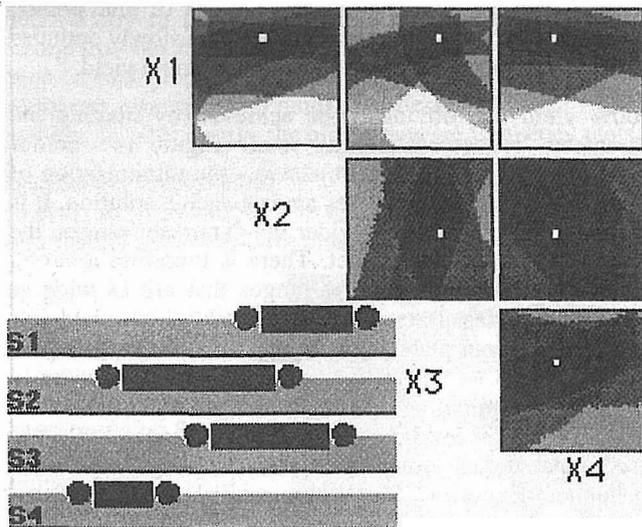


Figure 10: This Prosection Matrix represents 'slices' through parameter space. The grey scaling show how the data satisfies the performance requirements.

This is a projection of a section of parameter space, hence the name 'Prosection' (the term came from by a paper by Furnas and Buja [8]). This prosection process is repeated for every pair of parameters so that each scatterplot is displaying different data. The tolerance ranges for the scatterplots two parameters (p_1 and p_2 in figure 9) can also be projected on to the plot in the form of a tolerance box.

The Prosection Matrix shown in Figure 10 actually refers to a situation in which each parameter range is very small, leading to a small tolerance region (the small grey dot in the centre of each scatterplot). Because the parameter ranges are small, they define a very thin 'slice' through multi-dimensional parameter space, and therefore the resulting scatterplots show well-defined boundaries associated with the different performance limits of Figure 10. The colour coding used defines how well designs satisfy these performance limits. In Figure 10 designs that are acceptable are black, those that failed only one performance limit are dark grey and those that fail two are medium grey etc. One of the benefits of this colour coding is that the designer can explore the effect of moving the boundaries in the scatterplot. Thus, in Figure 11, the designer has moved the lower limit of performance S_3 even lower. A comparison of Figure 10 and 11 reveals how the corresponding boundary has moved, increasing the area of the (here, black) acceptable region. Exploration of this kind allows a designer to form a strategy for combining and trading off different performance requirements.

Though Figure 9 provides a conceptual illustration of the formation of each scatterplot within the Prosection Matrix it is actually unsuitable for implementation because it would result in a very grainy representation. Instead, each scatterplot is filled using a matrix of small coloured squares. For example if we consider the (top left) X_1X_2 scatterplot in Figure 11 its area is divided in 44^2 squares, The Cartesian coordinates of each square's midpoint defines values of X_1 and X_2 . Values of X_3 and X_4 are then

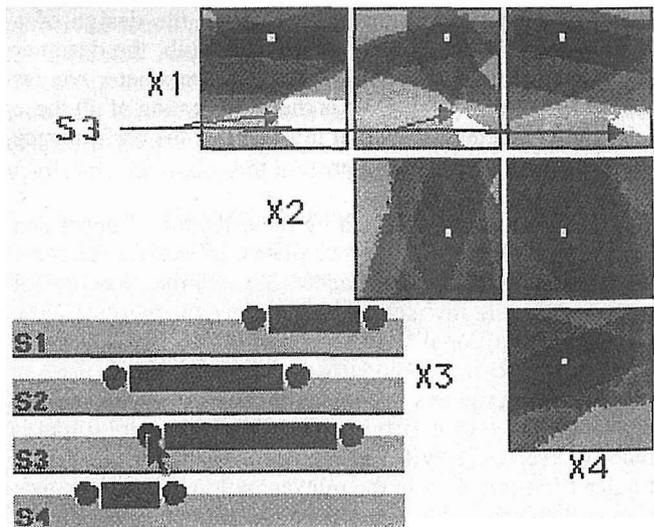


Figure 11: Adjusting a performance requirement and viewing how the related boundary moves in parameter space

selected randomly from within their tolerance range for each square. The corresponding values of the performances S1 to S4 are then computed from the model and compared with their respective limits. The square is then coloured according to the scheme already defined. For clarity, in the case of Figure 11 the X3 and X4 ranges are actually set to a single point so no randomisation occurs.

The existence of significant parameter *ranges* rather than single parameter values changes the detailed appearance of the Prosection Matrix but not its general character (Figure 12). Again consider the X1X2 scatterplot (top left). The original value of X1 has been replaced by a range of X1 as indicated by the yellow line. The immediate effect is that for all the scatterplots that don't have X1 as an axis, X1 is now randomly chosen within the selected *range* of X1 values rather than set at a single value. The increased fuzziness of these plots reflects this process. The rest of Figure 12 shows the effect of additionally assigning ranges to X2, X3 and X4.

Figure 13 (colour plate) shows how the Prosection Matrix looks when the performance and parameter limits are set as in Figure 7. The red regions now correspond to acceptable bulbs, whereas those that are manufactured lie within the yellow tolerance regions. The small percentage of red points within this region indicate a low yield (19%). In Figure 14 (colour plate) the user has set the tolerances to very narrow ranges to find a high yield (100%). Since wider

tolerance ranges are normally associated with lower cost, the designer will endeavour to make the yellow -bounded tolerance region as large as possible, perhaps even trading off manufacturing yield against cost. Figure 15 (colour plate) shows how the user has adjusted the parameter ranges so that they just fit inside the red region, resulting in much wider tolerances (potentially cheaper components) whilst maintaining a reasonably high yield (96%)

FORMATIVE EVALUATION STUDIES

The design of IVAs is difficult - it is often hard to judge what users will find intuitive and how an IVA will support a particular task. We have therefore carried out a number of formative evaluation studies at different stages of the IVA's development. Ten pairs of subjects were tested. They were all graduate engineers/scientists enrolled on PhD programs. The pairs worked together, first with the Influence Explorer, then the Prosection Matrix and finally both tools together. Reassuringly, each pair of subjects were able to complete a tolerance design task in about 30 minutes.

We learnt some very simple lessons from these evaluations:
 a) *Maximise the directness of the interactivity.* For example one version of the Prosection Matrix forced users to map their interaction from the sliders. However users preferred to select and drag the tolerance box directly.
 b) *Seek out the most crucial information and then represent it appropriately and simply.* The most obvious example of

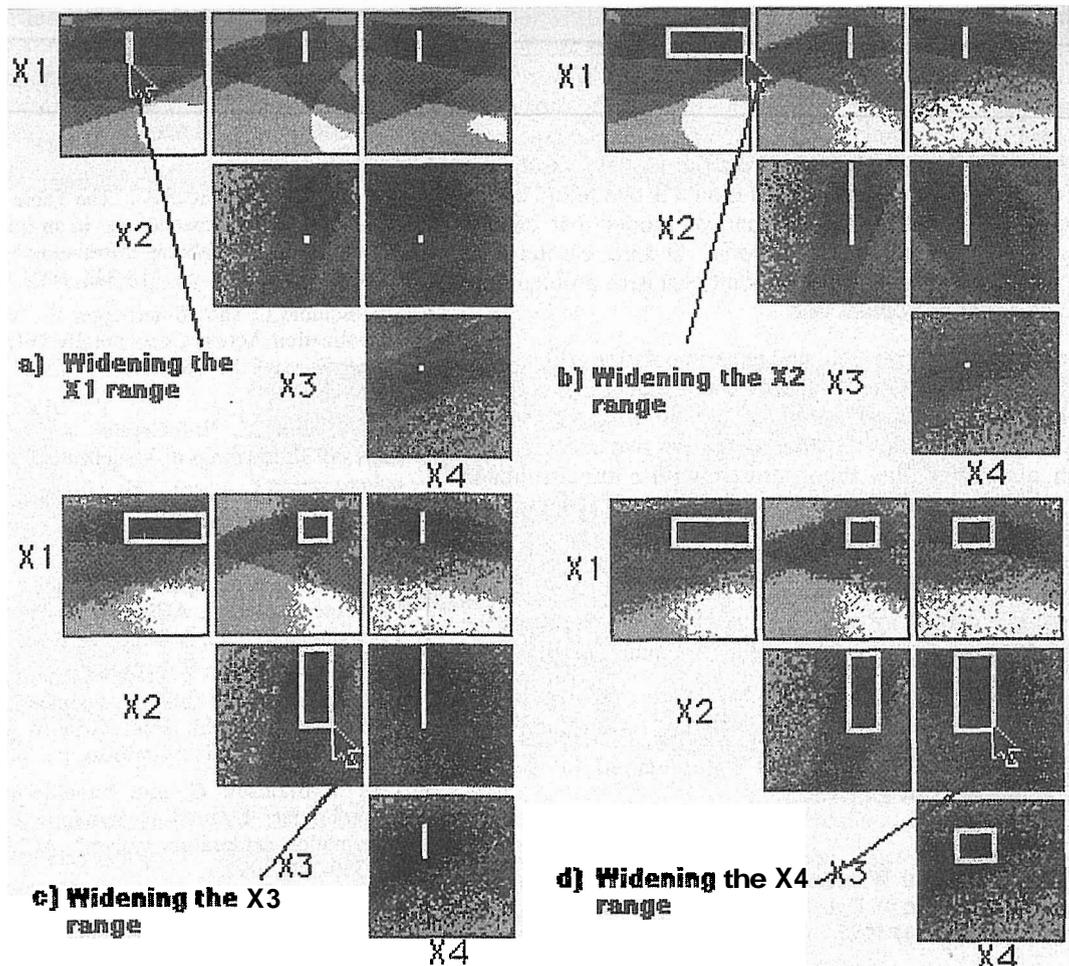


Figure 12: Gradually increasing the tolerance region so that sections of the data are projected. The boundaries become fuzzier as the ranges are adjusted.

this was the colour coding. Initially when considering the interface for setting up a performance specifications we attempted to colour code all the different variations of failure. Then we realised that this coding could be considerably simplified if we focused on encoding data that satisfied the performance limits and perhaps more importantly data that *almost* satisfied those limits. Colour coding the influence explorer for tolerance design was more difficult. The solution presented in this paper (Figure 8) has attempted to reduce the colour coded information to that which will provide immediate and useful information.

c) *There is a trade-off between the amount of information, simplicity and accuracy.* Ensuring that there is sufficient information to complete a task was an important issue. This emerged in the Influence Explorer when we tried adding tolerances with the original precalculated **dataset**. Using dynamic focused sampling overcame this problem.

CONCLUSIONS

The Influence Explorer and Prosection Matrix have now be utilised in a wide variety of industrial collaborations in electronic, structural and mechanotronic domains. The enthusiastic reaction of those who have observed and experimented with these **IVAs** suggests that the potential offered by immediately available and responsive interaction is considerable.

There are many reasons for this enthusiasm. One is the readiness with which opportunistic as well as planned exploration can be carried out. Another is the directness of external representations. Abstract Mathematical Models are difficult for the untrained user to interpret. However using these IVAs the problem holder can explore the model for themselves, and make use of their own considerable experience and knowledge to test the models validity in their own terms. A mathematical model is one thing, but an externalisation of that mathematical model that can be responsively explored is quite another. A third reason is that these tools transform a very difficult cognitive problem into a much easier perceptual task.

Many avenues of research and experimentation still need to be followed up. One concerns the enhancement of the designer's expertise by some of the automated tolerance design algorithms developed over the last two decades. One such algorithm was incorporated within the Influence Explorer and, when invoked, automatically and very rapidly (e.g. 10 seconds) adjusted the 'nominal value' of each parameter (the mid-point of the selected parameter range) to **maximise** the yield. Nevertheless, this automation needs to be complemented by an interface which will facilitate the human observation and guidance of automated design.

ACKNOWLEDGEMENTS

We thank Aarnout Brombacher (Philips, Eindhoven) and John Nelder. This work was implemented in C on a Macintosh by Huw Dawkes.

REFERENCES

[1] Ahlberg C. and Wistrand E., "IVEE: An Environment for Automatic Creation of Dynamic Queries Applications," **CHI'95** Demonstrations, May 1995.

- [2] Andrews D.F. "Plots of High-Dimensional Data" **Biometrics**, March 1972, pp 69-97
- [3] Becker R.A., Huber P.J., Cleveland W.S. and Wilks A.R., "Dynamic Graphics for Data Analysis", *Stat. Science* 2, 1987.
- [4] Bertin J., "Graphics and Graphic Information Processing", **deGruyter** Press, Berlin, 1977.
- [5] Chalmers M., "Using a Landscape to represent a corpus of documents", Springer-Verlag Proceedings of COSIT'93, **Elba**, pp. 377-390, September 1993.
- [6] Eick S.G., Steffen J.L. and Sumner E.E., " **SeeSoft™** - A Tool for Visualizing Line Oriented Software", *IEEE Transactions on Software Engineering*, pp. 11-18, 1992.
- [7] Feiner S. and Beshers C. "Worlds within Worlds: Metaphors for Exploring n-Dimensional Virtual Worlds", *ACM Proceedings 1990 Conference on User Interface Software Design*, pp 76-83
- [8] Furnas G.W. and Buja A., "Prosection Views: Dimensional Inference through Sections and Projections", *Journal of Computational and Graphic Statistics* 3 (4), pp. 323-353, 1994.
- [9] Inselberg A., "The plane with parallel co-ordinates", *The Visual Computer* 1, pp. 69-91, 1985.
- [10] Keim D.A. and Kriegel H., "VisDB: Database Exploration using Multidimensional Visualization", *IEEE Computer Graphics and Applications* September, pp. 40-49, 1994.
- [11] Nardi B.A. and Zamer C.L., "Beyond Models and Metaphors: Visual Formalisms in User Interface Design", *Journal of Visual Languages and Computing* 4, pp. 5- 33, 1993.
- [12] Newton C.M., " Graphics: from alpha to omega in data analysis", *Graphical Representation of Multivariate Data*, P.C.C. Wang (Ed) New York: Academic Press, pp. 59-92, 1978.
- [13] Mihalisin T., Gawlinski E., Timlin J. and Schwegler J., "Visualizing Scalar Field on an N-dimensional Lattice", *Proceedings of Visualization 90*, IEEE CS Press, pp. 255-262, 1990.
- [14] Rao R. and Card S.K., "The Table Lens: Merging Graphical and Symbolic Representations in an Interactive Focus + Context Visualization for Tabular Information", *Proceedings of CHI'94*, Boston, ACM Press, pp. 318-322, 1994.
- [15] Schmid C. and Hinterberger H., "Comparative Multivariate Visualization Across Conceptually Different Graphic Displays", *Proceedings of the SSDBM VII*, IEEE Computer Society Press, September 1994.
- [16] Spoerri A., "InfoCrystal: A visual tool for Information retrieval" *Proceedings of Visualization '93* pp150-157.
- [17] Suchman L.A., "Plans and Situated Actions - The Problem of Human-Machine Communication", Cambridge University Press, 1987.
- [18] Tweedie L.A., Spence R., Bhoghal R. and Williams D., "The Attribute Explorer", *ACM, Video Proceedings and Conference Companion, CHI'94*, pp. 435-436, April 1994.
- [19] Tweedie L.A., "Interactive Visualisation Artifacts: how can abstractions inform design?", *People and Computers X : Proc. of HCI'95* Huddersfield, (Eds) Kirby M.A.R., Dix A.J. and Finlay J.E., Cambridge University Press, pp. 247-265, 1995.
- [20] Williamson C. and Shneiderman B., "The Dynamic HomeFinder: Evaluating dynamic queries in a real estate information exploration system", *ACM, Proceedings SIGIR'92*, pp. 339-346, 1992.
- [21] Zhang J. and Norman D.A., "Representations in Distributed Cognitive Tasks", *Cognitive Science* 18, pp. 87-122, 1994.

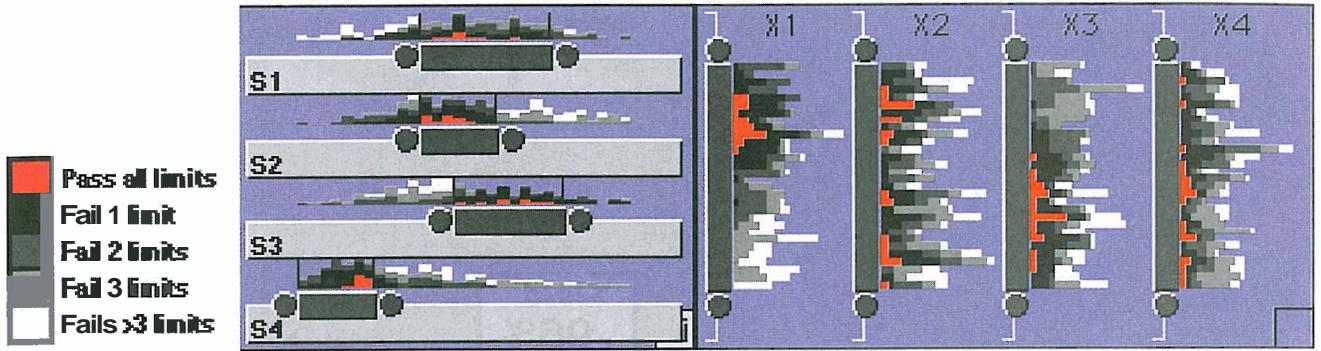


FIGURE 6

Setting up limits on the performance histograms (left).

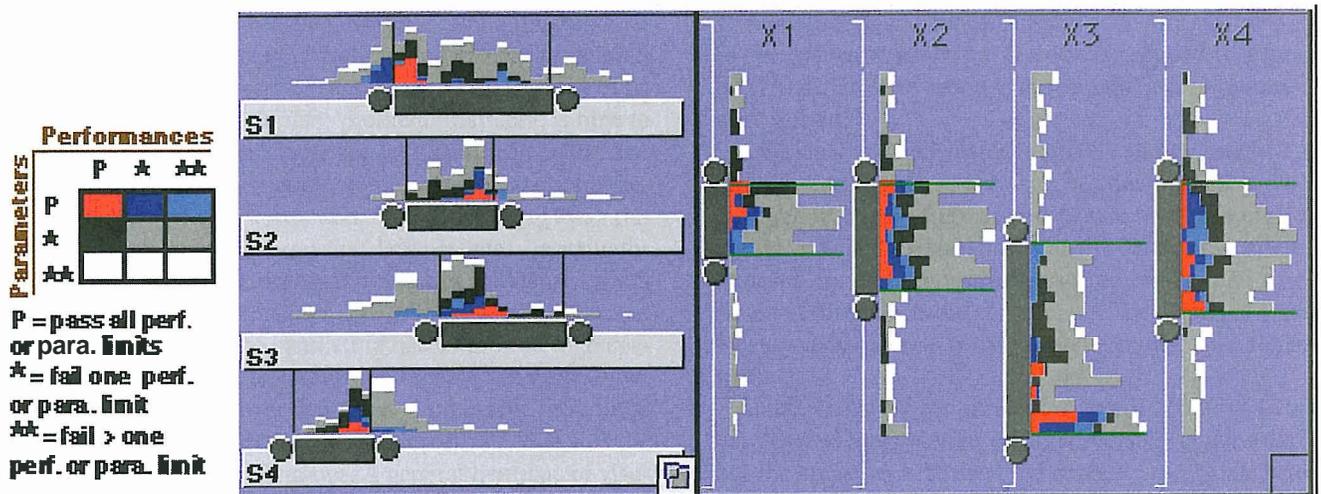


FIGURE 7

Setting up ranges on the parameter histograms.

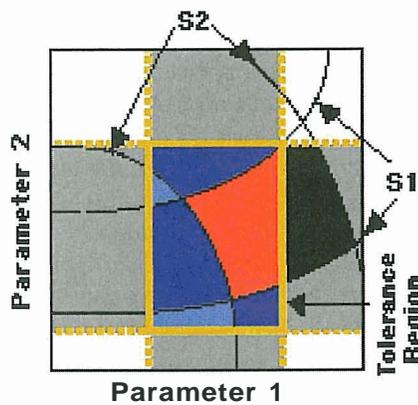


FIGURE 8

The colour coding for the tolerance setting (Figure 7) shown in 2 parameter space.

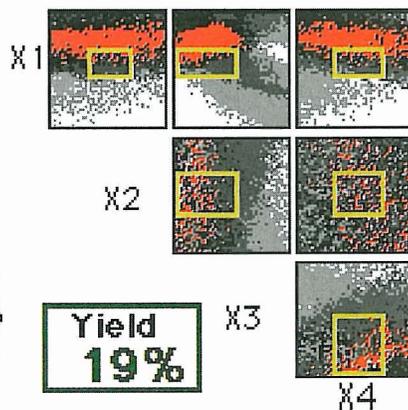


FIGURE 13

The Projection Matrix for the settings from Figure 7.

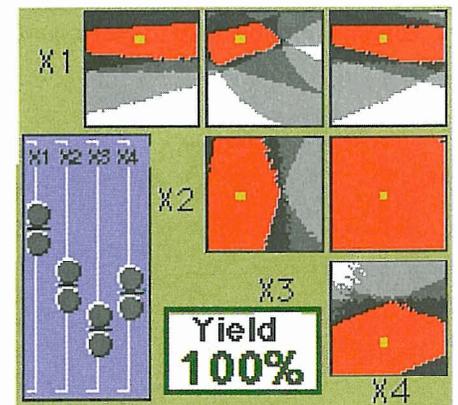


FIGURE 14

Here the parameter ranges are reduced to single values.

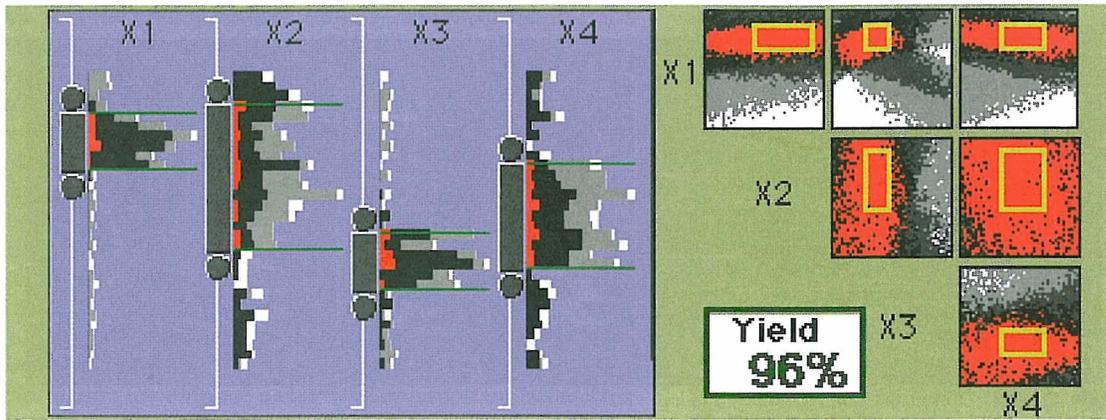


FIGURE 15

Maximising the widths of the parameter ranges to get high yield and minimum cost. (the same specification is used as in Figure 6). Note how the rectangular tolerance region just fits inside the red region of acceptability.