Visualisation For Functional Design

Bob Spence, Lisa Tweedie, Huw Dawkes, Hua Su Department of Electrical Engineering, Imperial College of Science, Technology and Medicine London, England SW7 2BT Tel: (+44) 171 594 6261 1.tweedie@ic.ac.uk

Abstract

In this paper we present two novel visualisation tools: the Influence Explorer and the Prosection Matrix. These were specifically created to support engineering artifact design and similar tasks in which a set of parameter values must be chosen to lead to acceptable artifact performance. These tools combine two concepts. One is the interactive and virtually immediate responsive display of data in a manner conducive to the acquisition of insight. The other, involving the precalculation of samples of artifact performance, facilitates smooth exploration and optimisation leading to a design decision. The anticipated benefits of these visualisation tools are illustrated by an example taken from electronic circuit design, in which full account must be taken of the uncertainties in parameter values arising from inevitable variations in the manufacturing process.

1 Introduction

The term "visualisation" has, unfortunately, virtually become synonymous with what has been called "scientific visualisation" and is associated primarily with 3D views of physical phenomena. Obviously such visualisation is highly beneficial in many circumstances. However, independent consideration needs to be devoted to significant problem domains in which 3D views are not naturally relevant in the physical sense. In such cases the associated data is largely abstract in nature, and the potential of visualisation is considerable and largely unexploited.

Engineering design is one such abstract domain. Its inherent difficulty and immense economic importance warrant, at the very least, an assessment of the potential benefits that visualisation can offer. In this paper we present a new visualisation technique appropriate to the needs of engineering design.

2 Engineering design

Whether an engineer is designing a toy, a silicon chip or a car, a primary objective in all cases, and one which is certainly shared with conventional 'scientific visualisation', is the acquisition of insight into some complex relationships. Additionally, however, a designer is required to act upon this insight and make specific choices of the constituent components of an artifact to ensure that its performance meets the requirements of a customer both reliably and economically.

It is characteristic of the design process that problem formulation, insight acquisition and problem solution proceed concurrently. It is also the case that the designer's concern may gradually progress from the qualitative to the quantitative. A further complicating factor is the dimensionality of the problem, aggravated by the fact that tens or hundreds of variables may be relevant to any design decision. Any visualisation tool proposed to facilitate the design process must satisfactorily address these and many other issues.

The **Influence Explorer** [1] and the **Prosection Matrix**, the subjects of this paper, may be such tools. To explain these anticipated benefits we must first examine, in outline, the process of engineering design.

2.1 Parameters and Performances

Central to engineering design is the relation between two types of variable. The **parameters** of an engineering artifact such as a bridge or electronic circuit are those variables whose value can freely be *chosen by the*





Figure 3: The iterative design process Right-to-left traversal (dashed) demands human expertise.

designer. For the electronic circuit designer the values of resistors and capacitors are designable; for the structural designer they may include the widths of girders and the depth of concrete foundations. Typically, the designer may have to choose as many as 50 parameter values. The values of these parameters are chosen to ensure, as far as possible, that the **performance** of the designed artifact satisfies the customer. Amplification and cost are representative examples of the performance of an electronic circuit, while structural performances might include the deflection of a bridge under maximum traffic load. As with parameters, the number of performances with which the designer must be concerned is typically large - 100 would not be regarded as unusual.

For simplicity we denote the set of designable parameters by $\mathbf{p} = p_1, p_2, p_3, \dots p_n$ and the set of performances by $\mathbf{F} = \mathbf{F}_1, \mathbf{F}_2, \mathbf{F}_3, \dots \mathbf{F}_m$. The nature of the relation between \mathbf{p} and \mathbf{F} is shown in Figure 1. A full description of the artifact - that is, a complete set of values for \mathbf{p} - uniquely determines the values of all the performances \mathbf{F} . Given \mathbf{p} , the corresponding values for \mathbf{F} can then be calculated by simulating the artifact.

The very reason that design is so difficult is that *there* is no direct inverse relation (Figure 2). In other words,

one cannot specify desirable values of \mathbf{F} , the performances, and directly and uniquely calculate the corresponding required set of parameter values \mathbf{p} . Indeed, some values of \mathbf{F} could easily be found for which a corresponding \mathbf{p} simply does not exist.

2.3 Conventional Design

In the absence of a direct transformation from \mathbf{F} to \mathbf{p} , the human designer will conventionally (Figure 3) propose, from experience or other sources of knowledge, a set of parameter values \mathbf{p} which is expected to lead towards acceptable performance. Simulation is then normally employed to determine the corresponding \mathbf{F} . If, as is frequently the case, \mathbf{F} does not fall within acceptable limits, the designer will change the \mathbf{p} values on the basis of experience. This iterative procedure ("trial and error") is repeated, often as many as 1000 times, until the specifications on performance are satisfied.

2.4 Human Expertise

As is evident from Figure 3, as well as from actual experience of design, the greatest impediment to design is the absence of any direct transformation between the performance specifications and appropriate values of the parameters. Human experience, supplemented by any insight that can be gained through exploration, is the only resource available and the design process is essentially iterative. The new approach we describe below continues to make use of human expertise (otherwise it would constitute a disastrously retrograde step), but places more emphasis upon smooth and continuous exploration and less on tedious iteration.

2.5 Exploration

To illustrate the new approach we consider the simple - though not trivially simple - case of two parameters and two performances (Figure 4). In **P**-space a point represents an artifact, and a point in **F**-space its (two) performances. First, a random selection of points in **P**-space is generated (Figure 4a) within a region called the **Region of Exploration**. The extent of this region is at the



Proceedings of the Proceedings on Information Visualization (INFOVIS '95) 0-8186-7201-3/95 \$10.00 © 1995 IEEE

discretion of the designer, and is usually anticipated to include an acceptable design. Typically the region of exploration will be quite wide, perhaps extending from 0.5 to 1.5 times the average value of each parameter.

Having generated a number of points in P-space (typically 100 or more), it is a straightforward matter* to simulate those artifacts and determine their performances (Figure 4b). In the simple case of two parameters and two performances the designer could then (Figure 4b, shaded region) select a point in F-space close to the desired performances and, thereby, immediately and without further calculation, identify the corresponding point in Pspace (Figure 4a, shaded region) and, hence, a reasonable approximation to a satisfactory design. Similarly, a choice of one or more available points in P-space will immediately indicate the corresponding points in F-space. It is the immediate exploration - in either direction, but especially from F to p - made possible by a reasonable pre-calculated coverage of *P*-space and *F*-space that lies at the heart of the two tools we will describe here.

3 The Influence Explorer

Application of the new approach - the precalculation of many \mathbf{p}/\mathbf{F} pairs - to a practical design problem obviously involves many more than the two parameters and performances illustrated in Figure 4. The consequent difficult but fundamental problem of displaying multidimensional data effectively on a two-dimensional display has received increasing attention recently [2]. The approach we have adopted is illustrated in the figures below which show the various 'windows' which comprise the interface of our first abstract visualisation tool, the Influence Explorer.

This tool makes use of the "Interactive Histograms" technique. This technique was an idea that came to us after reading Williamson and Shneiderman's work [3] where they used sliders as input mechanisms to visualisations. It became apparent to us that such sliders could also be used to display output. We first tried this in "The Attribute Explorer"[4], a visualisation tool for searching for a house within a large population of houses. The Influence Explorer takes the Interactive Histograms technique a stage further by applying it to a more complex problem.

In the following section we will describe the functionality of the Influence Explorer. We first describe the Performances Window and the Parameters Window:

Performances Window The performance of any given artifact is represented by a rectangle appropriately placed on each of the performance scales (Figure 5). Thus, the collection of samples generated in the manner shown in Figure 4 is represented by a histogram on each of the performance scales, as shown in Figure 5; each histogram is composed of the same number of rectangles as there are sample points in the Region of Exploration. The designer can place upper and lower limits, indicated by vertical bars, on each performance, whereupon the artifacts falling within all those limits are identified by colour (red) on all the performance scales. The illustration refers to actual

data concerning electronic circuit design, as will be explained later.

Parameters Window In terms of parameter values, the description of any artifact or collection of artifacts is represented in the same way as for performances (Figure 6). Again, there is a facility, explained later, for placing limits on parameter values. A single artifact is therefore represented once on each parameter scale and once on each performance scale.

3.1 Parameter/Performance Relations

As an illustrative example we will use the design of an electronic amplifier circuit [5]. This is described by five parameters related to the dimensions of various components in the circuit and four performances exhibited by the amplifier. Figures 5 and 6 refer, in fact, to this example and display the result of simulating the performance of 800 circuits in which the parameter values have been chosen within wide exploratory ranges.

When the histograms first appear in the performance and parameter windows the designer may well wish to gain some feeling for numerical values and the location of specific circuits on each of the scales. Selecting a rectangle on any scale flashes corresponding rectangles on all the other scales and displays all the numerical values (Figure 5). A useful visualisation mechanism is also provided if the same circuit is joined by lines on each scale (Figure 5): collections of such lines can provide useful insight into the relationships implicit in the displayed data. These lines are known as parallel co-ordinate plots [6,7].

3.2 Early Exploration

It is a characteristic of engineering design that, in its early stages, the designer tends to think qualitatively. Such an activity is well supported by the Influence Explorer. For example, a window or range can be selected on one scale. Immediately, all the rectangles corresponding to the selected circuits will be coded red. The designer may, in fact, thereby define a small 'window' on a performance scale and then, by mouse action, move that window backwards and forwards along the scale: the result may well be a well-defined movement of red squares on another performance scale, indicating a correlation or trade-off, knowledge of which is immensely helpful to the designer. At the same time, the movement of the red squares on the parameter scales may indicate a strong influence of one or more parameters on one or more performances.

Exploration of the overall relations implicit in the histograms can be facilitated by the display of yellow circles (Figure 5) which indicate, for each scale, the

^{*} How straightforward and costly depends upon the complexity of the model. For the illustrative example employed in this paper a Response Surface Model - comprising a few simple and readily executed expressions - was obtained for the artifact by fitting the model to a reasonably small number of accurate simulations [3].

average value of the red squares: this display is activated from the tools window.

3.3 Specifications on Performance

At some point the designer will wish to begin taking into account the existence of limits to various performances. As already explained, activation of the appropriate tool allows upper and lower limits to be placed on the performance scales (Figure 7), whereupon all circuits that satisfy all the limits are colour-coded red. For the remaining circuits it is also useful to employ colour coding to indicate how many performance limits are violated: black indicates the violation of only one limit, dark grey items fail two limits and so on.

In Figure 7 we see that there are only a few red points, though there are many black points (failing only one limit) just beyond the right edge of the upper limit on DC Gain. By adjusting this upper limit to the right, therefore, many black points will become red. The designer immediately forms some idea of the severity of the DC Gain upper limit as far as the number of satisfactory circuits is concerned.

Examination of the location of red squares on the parameter scales will indicate where possible designs might lie, and will often help to identify the relation between performances and parameters. With the performance limits in place, any red square represents a circuit which satisfies all performance requirements.

3.4 Effects of Uncertainties

It is important to note that a performance limit will be introduced for one of two reasons. In the previous section the limits were introduced to explore their effect on the number of red squares in existence and the locations of acceptable parameter values. Later in the design process, however, consideration must be paid to the requirements of the customer. Where these requirements are negotiable, the exploration of the previous section can provide valuable information.

Where the performance limits are rigid, the artifact designer must take account of the fact that, with a massproduced product, inevitable variations in the manufacturing process lead to uncertainty regarding parameter values. As a consequence, the value of each parameter can only be guaranteed to lie somewhere within a given *range*. Clearly, it is desirable that the fraction of manufactured artifacts that satisfy the performance limits be as high as possible. The fraction that does so is called the **manufacturing yield**.

The problem faced by the designer can be illustrated by reference to a two-parameter artifact (Figure 8). The ranges within which the two parameters are known to lie describe a rectangle (shaded), called the Tolerance region, in parameter space. The limits identified by the customer in performance space transform to generally very complex surfaces in parameter space. Together, these surfaces define a **Region of Acceptability**. Thus, the artifacts that are manufactured lie within the Tolerance Region, but only



Figure 8: The Tolerance Design Problem

those which also lie within the Region of Acceptability will satisfy all the performance limits. For a real artifact with many performance requirements the same concepts illustrated in Figure 8 apply, but in a very much higherdimensional space.

3.5 Tolerance Design

The designer will consider two objectives. The first is to maximise the manufacturing yield, which in the simplest case is that fraction of the Tolerance Region lying with the Region of Acceptability. To do this the location of the Tolerance Region (Figure 8) must be adjusted suitably within parameter space. Having maximised the yield, the designer must try to minimise the cost of mass-producing the artifact. The relevant fact here is that, for most components within an artifact, a wider tolerance range is usually associated with a lower cost. The designer will therefore try to inflate the Tolerance Region of Figure 8 as much as possible without causing a serious fall in yield: the eventual satisfactory compromise will depend upon local knowledge of economic factors.

The act of maximising the yield, as described above, is called Design Centering [8]. That of inflating (or deflating) the size of the tolerance region is called Tolerance Assignment. Together, these activities are referred to as Tolerance Design. The reason why tolerance design is difficult is that the location of the Region of Acceptability is unknown and its computation prohibitively expensive.

Tolerance Design can be achieved either manually by the designer or automatically by means of an optimisation algorithm. Both approaches are considerably facilitated by visualisation tools.

3.6 Manual Design

The expected range of a parameter value can be specified in precisely the same way as for a performance limit. The colour coding is now supplemented by blue squares on the parameter histograms as shown in Figure 9 where the tolerance ranges are set to their maxima prior to possible reduction. Each blue square corresponds to a circuit whose parameter values all lie within the parameter tolerance ranges but whose performance is unsatisfactory: these ("fail") circuits are the ones that, after manufacture, would be found unsatisfactory. The ratio of red ("pass") circuits to the total of red and blue circuits is called the manufacturing yield.

To achieve design centering the positions of the parameter ranges are manually adjusted. Immediately, and with little delay, (thereby encouraging exploration), the yield is calculated and displayed in the Tools window.

4 The Prosection Matrix

Following design centering, tolerance assignment may be attempted. This aspect of design is facilitated by the second tool called the Prosection Matrix appearing in the Prosection Window (Figure 10). This is illustrated in Figure 11 for the example of a three-parameter artifact. The rectangular areas (numbering 5!/2 for the 5-parameter example in Figure 10) within this tool correspond to all possible pairs of designable parameters. The small dots each corresponding to a different circuit - within each area are coloured red and black according to whether the circuits they represent pass or fail the performance specifications. The position of each square is determined solely by the values of the two parameters, but only those points corresponding to circuits defined by the remaining parameter ranges are displayed. Indeed, the Prosection Matrix is so-called because it shows a projection, onto a plane defined by two selected parameters, of all points in P-space associated with that section of P-space defined by the choice of parameter ranges. The mathematical basis of combining sections and projections has been discussed by Furnas and Buja [9].

Also displayed in each area of the Prosection Matrix is a 'tolerance box' corresponding to the tolerance range for each of the two parameters involved. Examination of the relative location of the green tolerance box and the red (passing) and black (failing) circuits enables a judgement to be made as to whether it may be possible to enlarge one or more tolerances in order to reduce the manufacturing cost without seriously affecting the yield. Exploration



Figure 11: A section of p3 is projected onto a p1/p2 scatterplot

based on this judgement is assisted by the continuous display of the yield estimate and the visibility of the parameter tolerance ranges.

5 Automated Design

As already mentioned, both design centering and tolerance assignment can be achieved automatically by appropriate algorithms [8]. Nevertheless, a display of their progress and final result in both these tools achieves the beneficial result of complementing the power and speed of algorithmic procedures with the insight, interpretation and experience of the human designer.

6 Comments and Conclusions

Experience already gained with these tools in the field of structural and electronic design, as well as the comment of industrial designers, suggests that visualisation tools have much to offer in fields characterised by abstract rather than 3D representations. Specifically we feel that the tools presented here may well point the way towards a new generation of design tools.

Many aspects of the Influence Explorer and Prosection Matrix remain to be investigated and developed. For example, the ability of the tools to integrate the value of algorithmic procedures and human expertise opens up a huge field of investigation which has received little attention [10]. We hope to carry out some evaluation of these tools although the complex nature of the design task may make this difficult.

7 Acknowledgements

We thank John Nelder and Aarnout Brombacher. This work is funded in part by ESPRC UK and Philips Research.

References

[1] L.A. Tweedie, R. Spence, H. Dawkes and H. Su *The Influence Explorer* Interactive Poster, Conference Companion CHI'95 pp 129-130 April 1995.

[2] W.S. Cleveland and M.E. McGill "Dynamic Graphics for Statistics" Wadsworth and Brooks/Cole, Pacific Grove, CA. 1988.

[3]C. Williamson and B. Shneiderman "The Dynamic HomeFinder: Evaluating Dynamic Queries in a Real-Estate Information Exploration System", SIGIR'92, Special issue of the SIGIR Forum, Nicholas Belkin, Peter Ingwersen, and Annelise Mark Pejtersen (Eds) Copenhagen, Denmark, pp 338-346, June 1992

[4] L.A.Tweedie, B.Spence, D.Williams and R.Bhogal "The Attribute Explorer" ACM, Video Proceedings and Conference Companion CHI'94 pp435-436 April 1994

[5] H. Su, C. Michael, and C. Ismail "Statistical Constrained Optimization of Analog MOS circuits using Empirical Performance Models" IEEE, Proc. ISCAS'94, pp. 133-136, May 1994, London.

[6] A. Inselberg *"The plane with parallel co-ordinates"* The Visual Computer 1, pp. 69-91, 1985.

[7] E.J. Wegman, "Hyperdimensional Data Analysis using Parallel Co-ordinates" Journal of the American Statistical Association, Vol 85 no. 411, Theory and Methods, pp. 664 - 675, 1990.

[8] R. Spence and R.S. Soin "Tolerance Design of Electronic Circuits" Addison -Wesley 1988.

[9] G.W. Furnas and A. Buja "Prosection Views: Dimensional Inference through Sections and Projections", Journal of Computational and Graphical Statistics, 3(4), 1994, pp. 323-353. (American Statistical Assoc).

[10] L. Colgan, P. Rankin and R. Spence "Steering Automated Design" Proceedings of the Conference on Artificial Intelligence in Design '91 Oxford, Butterworth-Heinemann, pp 211-230, 1991

Proceedings of the Proceedings on Information Visualization (INFOVIS '95) 0-8186-7201-3/95 \$10.00 © 1995 IEEE

-a. Figure 5: "Performances" window Figure 6: "Parameters" window



Figure 7: Selection of performance specifications



Figure 8: Selection of tolerance specifications

