

# Navigating High-dimensional Spaces to Support Design Steering

Helen Wright<sup>1</sup>

Ken Brodlie<sup>2</sup>

Tim David<sup>3</sup>

Department of Computer Science,  
University of Hull, UK

School of Computer Studies,  
University of Leeds, UK

Department of Mechanical  
Engineering, University of Leeds, UK

## Abstract

Throughout the design cycle, visualization, whether a sketch scribbled on the back of a spare piece of paper or a fully detailed drawing, has been the mainstay of design: we need to see the product. One of the most important stages of the design cycle is the initial, or concept, stage and it is here that design variants occur in large numbers to be vetted quickly. At this initial stage the human element – the designer – is crucial to the success of the product.

In this paper we describe an interactive environment for concept design which recognises the needs of the designer, not only to see the product and make rapid modifications, but also to monitor the progress of their design towards some preferred solution. This leads to the notion of a design parameter space, typically high-dimensional, which must also be visualized in addition to the product itself. Using a module developed for IRIS Explorer<sup>TM</sup>, design steering is presented as a navigation of this space in order to search for optimal designs, either manually or by local optimisation.

**CR Categories:** I.6.7 [Computing Methodologies]: Simulation and Modeling – Simulation Support Systems; J.6 [Computer Applications]: Computer Aided Engineering – Computer Aided Design

**Additional Keywords:** Computational steering, design steering, concept design, multidimensional visualization, scientific data visualization.

## 1 INTRODUCTION

Visualization as a data analysis tool plays an important rôle in computational science. Although often applied in a post-processing step after the simulation is completed, a number of workers have noted the benefits of a more integrated approach [see for example 6, 7, 8, 9]. Called *computational steering*, parameters of the simulation can be altered whilst it is

in progress and the results re-calculated and visualized at once, leading to a more rapid solution of the problem in hand.

In this paper we examine the scope for computational steering in the particular domain of product design, leading to an approach which we call *design steering*. Although related to computational steering in broad terms, we see design steering as being distinctive in two key respects. Firstly, the crucial variables of computational steering and design steering are different. The former deals with simulation parameters which might include, for example, the mesh geometry, the error tolerance, or the time step to use. The latter deals with design parameters such as product geometry, mass and cost, and performance measures such as efficiency and failure rate. Secondly, the goal of computational steering is in some sense to validate our understanding of a physical process, that is, to ensure our computational model reproduces behaviour observed in the real world. Although correct behaviour is also a requirement for design, the primary aim of design steering is to optimise the performance of the product to which this model is then applied. Thus, if we adopt Johnson's analogy [7] of using computational steering to debug the simulation process, we can view design steering as an outer loop which takes place once we are confident that our computational model is correct.

A number of studies have addressed the informational needs arising from computational steering. One approach put forward by Brodlie *et al.* [3] was to model the investigation process as a history tree, where branch points of the tree corresponded to changes made to parameter values. More than just computational steering, where changes are made only at the current point reached by the simulation, the history tree allows the scientist to backtrack to some earlier stage before continuing with a modified set of parameter values. This work later found expression in the visualization package IRIS Explorer<sup>TM</sup> [16], as an example of support for computational steering within dataflow systems. Another approach proposed by Mulder and van Wijk is to log the history of the data as it changes during the simulation [10], by linking the appearances of graphical objects to parameter values in the simulation.

The informational needs of design steering are somewhat different, however. Here the requirement is akin to a decision support system, not only recording past design attempts, but also guiding the designer towards improved products. We begin by examining the nature of design problems and propose a visualization metaphor wherein individual design attempts are used to populate a high-dimensional space accommodating both the design variables and quality criteria. By assigning quality values to individual design attempts we can then postulate new designs, either automatically by applying optimisation techniques or by virtue of the human designer's own experience of the product being modelled. The paper then goes on to describe our presentational approach for this information.

---

<sup>1</sup> h.wright@dcs.hull.ac.uk <sup>2</sup> kwb@scs.leeds.ac.uk <sup>3</sup> t.david@leeds.ac.uk

## 2 CONCEPT DESIGN

One of the most important stages of the design cycle is the initial, or concept, stage where the need is to vet a large number of design variants before choosing to subject a small number to further analysis. Here the ideas should flow, neither hindered by lengthy computational analysis nor bogged down in parameter values and complex drawings. At this point in the design cycle the designer, the human element, is vital to the success of the product, contributing knowledge and experience which is not easily captured in terms of hard-and-fast design principles.

To illustrate, consider the design of a radiation furnace, of the type commonly found in the glass-making industry, whose configuration is such that a concave roof and rectangular walls and floor provide an enclosure where heat is generated by combustion so as to melt the glass. The process of glass manufacture is continuous and the design problem in this case is to ensure that 'hot spots' and large temperature gradients do not occur in the furnace walls, since this degrades the refractory material and leads to expensive shut-downs. The burning gas in the chamber usually has the highest temperature in the system and is a known constant, however, the temperature within the walls and its variation must be solved using a coupled boundary element method [12].

Since the boundary element method requires only knowledge of the domain boundary, and matrix elements are formed by virtue of the geometric interaction of pairs of elements, then the matrix is a function of the enclosure geometry. In order to provide a rapid evaluation we use efficient matrix storage mechanisms and relevant matrix elements are only updated when changes are made to an associated boundary.

Once the boundary value problem has been solved then the temperature can be found at any given position within the furnace wall. In order to obtain a reasonably small set of design parameters the roof geometry is modelled as the solution to a 4th order partial differential equation where the design parameters are mapped to the vector boundary conditions, i.e. the tangent vector of the curve at the boundary. The placement of the enclosure within the furnace walls may also be varied. Figure 1 shows the basic design of the furnace, together with these parameters.

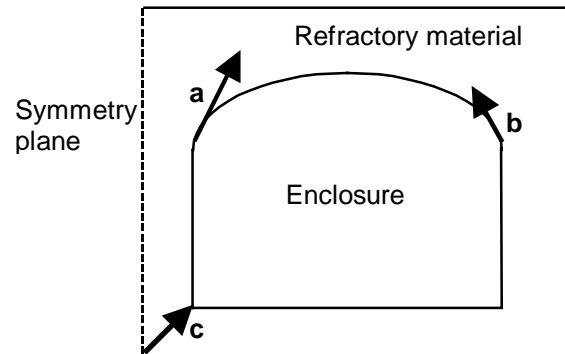
The design cycle in this case can thus be described as

1. new geometries are formed by alteration of the six design 'handles' ( $x$  and  $y$  components of each of the three vector quantities  $a$ ,  $b$  and  $c$  denoted in Figure 1)
2. the boundary problem is solved, and the temperature evaluated throughout the solid
3. contours of temperature and temperature gradient are available for inspection and the geometry is modified on the basis of this and knowledge of previous attempts


To support this cycle we have built a system using IRIS Explorer<sup>(TM)</sup>, whereby the designer can alter designs and calculate the temperature characteristics as described above, so as to visualize the results and make decisions quickly. The system in operation is shown in Figure 2.

In its realisation, the concept design environment for radiation furnaces described here is similar to that for prosthetic heart valves proposed by David and Hsu [4], where a quick design-and-analyse approach is adopted to modelling the flow of blood

through the opening and closing valve. Although very different applications, the aim of the approach in each case is to minimise the design-to-product cycle time by reducing the number of designs for which a full, complex and time-consuming computational analysis is required. Speed is of the essence and, coupled with the specification of product geometry by means of a small number of design handles, we believe this design philosophy represents the way forward for a wide variety of products.



**Figure 1** Radiation furnace comprising furnace walls and an enclosure, with vectors to denote the designable parameters of the system. Furnaces occur in pairs so the left-hand opening is modelled by symmetry. In a continuous process, melting glass is conveyed as if into the page, thus by neglecting end effects the problem is solved in cross-section.



**Figure 2** Concept design environment for the radiation glass furnace application. All of the parameters shown contribute to the specification of the furnace geometry, but particular attention is paid to those labelled. Altering any slider or dial immediately results in re-calculation of the temperature contours.

### 3 DESIGN STEERING

An integrated system such as that in Figure 2 shows the obvious parallels between computational steering and design steering, and clearly visualization plays a part in assessing designs. In design, however, this is not the only rôle for visualization, since there is another space that can be represented in addition to the physical space of the product and its surroundings. This is the space containing the inputs to and outputs from the design process itself; in the case of radiation furnace design these respectively would be parameters defining the shape and position of the enclosure and critical values produced by the subsequent analysis, such as mean and maximum temperature in the solid domain, and mean and maximum gradient of temperature likewise.

#### 3.1 Relationship to Prior Work

Tweedie *et al* [15] aptly describe this space as the *abstract domain of engineering design*. In their description of artifact manufacture they draw attention to the use of simulation to go from a set of parameters to a set of performances, and also succinctly state the problem of design as the lack of any inverse relation between these two. Their approach is to regard the design domain as two related spaces, one a parameter space and the other a performance space. The dimension,  $N_p$ , of the parameter space will be determined by the number of parameters influencing the design and the dimension,  $N_s$ , of the performance space by the number of performances calculated. Thus the coordinates of a point in parameter space will depict a particular design and those of the corresponding point in the related space its performances. Tweedie *et al* populate their design domain with precalculated parameter-performance pairs, in order to select a range of parameters that define a set of designs exhibiting the required performance(s). Using a variation of the multidimensional scatterplot technique, visualization of these selected designs is then accomplished by projecting sections of the space onto the  $N_p(N_p-1)/2$  two-dimensional planes formed by all possible distinct pairs of parameters. Design variants in this so-called *prosection matrix* are points coloured according to whether they pass or fail certain design criteria.

Shaffer *et al* [14] and Goel *et al* [5] in their studies of aircraft design also recognise the potential for visualizing design variables, the former slicing through parameter space to see a subset of designs and the latter employing a technique called parallel coordinates in order potentially to see all design attempts at once. Both also provide a view of design quality, by colouring the slice containing the design points or the individual polygonal lines which comprise the parallel coordinate display. From the point of view of steering, a distinctive feature compared with Tweedie's work is the ability to specify designs as they are required, by means of sliders on a user interface, rather than by pre-calculating them. These new designs are added to any already present in the database, and the set visualized as a whole.

#### 3.2 Design Space as an Information Terrain

Our approach has likewise been to combine the parameter and performance spaces into one representation, but to present this in such a way that the designer can steer from within it, as well as

using it to review designs and their relationships. Using the space to specify designs obviates the need to switch between different tools on the user interface and helps to maintain the context provided by earlier design attempts.

We refer to this design steering process as *navigation* because we envisage the designer moving around the abstract design space much as they would move around the real space we inhabit, pausing at various points to specify and simulate designs 'on-the-fly', and leaving a permanent record of their traverse for subsequent review. The approach we have taken is to treat the design space as an information terrain, in the manner proposed by Benford *et al* [2] in their study of the support needed for cooperative work in virtual environments. They describe an information terrain as an abstract space within which objects representing the data can be displayed and manipulated. Benford and co-workers demonstrated their idea in support of the representation of entries in a database. In our work the purpose of the terrain is to demonstrate relationships between the design variants by means of the spatial distribution and appearance of the objects representing them.

#### 3.3 Extrinsic and Intrinsic Dimensions

In its abstract conception the dimension of the design space is the sum of the number of parameters,  $N_p$ , and performance measures,  $N_s$ , of the product we are simulating. If we are to present this information to the designer visually we need a means of mapping the  $N_p + N_s$  pieces of information such that they are easily perceived and manipulated. Benedikt [1] introduces the concept of extrinsic and intrinsic dimensions of point-objects in Euclidean space, a concept which Benford and co-workers employ to good effect in their information terrain. Extrinsic dimensions comprise the location of the point-object in space and time, whilst intrinsic dimensions describe its colour, size, shape and so on. For example, a cylinder placed at a point on a 2D plane might represent two design parameters by means of its  $x$  and  $y$  coordinates, two further design parameters by using its length and radius, and a performance measure by employing the saturation value of its chosen colour.

This method of representing designs presents two principal difficulties, however. Firstly, the number of extrinsic dimensions that we can utilise is limited to the three spatial dimensions we can ordinarily perceive, such that additional design parameters have to be accommodated using intrinsic dimensions. If we are to use the space in order to specify a design, then the designer must not only stipulate the coordinates of an object but also its physical characteristics – a task that is both complex to support in 3D graphical interaction, and likely to lead to mis-specified designs. In short, the small number of extrinsic dimensions available to us complicates the sense of 'going' to a point in parameter space to execute a design. Secondly, when the space is used to review performances, the use of intrinsic dimensions for both parameter and performance variables alters the designer's perception of a viable solution, since the tendency is to sum the obvious characteristics of an object, such as its size, height and colour saturation, into an overall indication of merit, notwithstanding that some of these variables may have nothing to do with quality.

Both difficulties point to an approach such that extrinsic dimensions always describe parameter values, and intrinsic dimensions always depict performance measures. Typically there is no shortage of intrinsic dimensions – object size, colour

saturation and hue are all viable for depicting quality information, but the challenge is to increase the number of extrinsic dimensions sufficiently. One technique, echoing Tweedie's work, would be to employ a multidimensional scatter plot, projecting the points representing the design variants onto a tiling of  $Np(Np-1)/2$  two-dimensional planes. A problem with this approach in the context of design steering is immediately apparent though, namely, the need to associate points on different planes with any one design variant, so that the designer knows which one is currently under scrutiny.


### 3.4 Implying 'Presence' in the Design Space

A number of solutions to this difficulty can be proposed but one in particular, previously used to visualize chemical reactions [17], turns out to be especially useful in reinforcing the sense of navigating through the design space. To illustrate, suppose there are six parameters yielding fifteen 2D scatterplots, from which three are chosen such that each parameter occurs exactly once. Each design is represented by a triplet of points and the coordinates of each triplet, six in total, serve to denote the parameter values. These selected scatterplots are then drawn on orthogonal planes – the planes' intersection serves as the origin of the six-dimensional coordinate system, translated to the point at which the first design is specified. It is then a simple matter to join the points together in the order they are created, forming a trajectory through the six-dimensional space. The most recent design variant is always at the head of the trajectory, whose three arms (effectively its two-dimensional shadows) emanate from the first design placed at the origin. There remains the question of how to associate together points in the middle portion of the design path, but in practice we find this is only a problem with very long trajectories. In this event, creating a new display from the current point is usually beneficial. From the designer's point of view, steering the product's design now becomes a walk through parameter space, leaving behind a trail of design variants.

To support this navigation process we have implemented a module for IRIS Explorer<sup>(TM)</sup>, called **SpaceWalk**, which is compatible with any design steering application such as in Figure 2. Since any parameter can go up or down from its initial value, we can if necessary use semi-transparent planes on which to draw the three arms, so that all the elements of the trajectory can be seen from any orientation. If the design space is of higher dimension than six, additional parameters can be accommodated in multiples of this basic space, switching between them by means of a slider on SpaceWalk's user interface. Figure 3 shows the SpaceWalk geometry produced for the radiation furnace design described above, with a performance measure based on both temperature and temperature gradient being used to control the colour value of the points drawn. A non-linear mapping of value to performance measure is used, such that only a small proportion of the best designs show up as lighter coloured points. Other mappings that could be used to show multiple performance measures simultaneously include the colour saturation, transparency and, possibly, the radius of the points.

Specification of new designs, and the review of previous ones, is done by direct interaction with SpaceWalk's geometry. This allows the user to feel present within the design space. For example, to specify a new parameter set, the user clicks on each of the three orthogonal planes at the point where they wish to see

the next design, with the choice of position guided by the presence of previous attempts. SpaceWalk converts this information to numerical values and passes these to the design geometry generator. The boundary problem is solved and the temperature within the solid is calculated in order to display the temperature contours, and a performance measure is generated indicating the quality of this design. This information is returned to SpaceWalk to draw the design at the appropriate point (the new head of the trajectory) and colour it. If the new attempt is worse than some earlier one, the previous best (or, indeed, any) design can be retrieved by clicking on one of the three points representing it. Retrieved designs are highlighted in the space, which further emphasises the association of a triplet of points into a single design variant.



**Figure 3** A short walk through six-dimensional design space. Apart from the first design lying at the origin, a triplet of points is needed to record each design attempt. The light colour of the designer's second attempt shows it is the best; this can be recalled by clicking on one of the points representing it. At the same time as the design space is visualized, so are the furnace geometry and temperature contours.

## 4 OPTIMISING DESIGNS

Implicit in what we have described above is a sense that the design process is an optimisation activity, regardless of whether this is expressible in any formal way. Whereas in computational steering the aim is to make the simulated system mimic some observed behaviour, the designer, on the other hand, will have in mind a set of criteria against which to measure the quality of their product. Some of these  $N_S$  performances might be inferred visually, whilst others might be calculated, but in terms of design activity the goal is the same – to learn from all the variants tried so far in order to steer the design towards an improved product.

### 4.1 Manual Search

We can see this learning effect in a manual search such as that shown in Figure 3, where by chance the second design tried is of fair quality. Hoping this heralds the start of a pocket of good designs, the designer proposes a third a short distance further on

in the space, keeping the roof angle the same at the right-hand edge but making it more vertical at the left-hand one. However, the combined effect of these changes is to make the roof more concave and, coupled with the shift right of the whole enclosure, they markedly reduce the performance measure. In terms of its location in the space, this new design overshoots the fruitful region. After a further attempt in the same part of the space to confirm this diagnosis, the designer returns to a point near to the earlier, better design.

## 4.2 Automatic Search

Whether or not the search process can be automated is dependent on the type of performance measure that can be defined. If a design's performance can be calculated numerically at the end of each iteration of the design cycle, it is possible to apply numerical optimisation techniques to maximise the performance with respect to the design parameters.

We have developed an optimiser for the radiation furnace application based on the downhill simplex method of Nelder and Mead [11], and following the algorithm suggested by Press *et al* [13]. This method was chosen because it is robust and does not require the calculation of derivatives. The method is started by giving it an  $Np$ -dimensional point as one vertex of the simplex; the other  $Np$  vertices are calculated so as to span the available ranges of the parameters. The function minimised reflects how distant the current design is from some ideal performance set by the designer. This is calculated at each vertex to find the current lowest and highest, and the optimisation proceeds via a series of reflections, expansions and contractions of the simplex, aimed eventually at drawing all  $Np+1$  vertices together around a single low point. The algorithm *does* only find a local minimum but, as we shall see, this does not diminish its usefulness.


## 4.3 Combined Search Approach

An important feature of our system is the possibility to combine a manual and automatic search as needed. Thus we find the optimiser very useful for honing designs produced manually. In Figure 3 the last design attempt was used to start the optimiser and in a short space of time it produced a better design, but with broadly the same geometrical characteristics as the designer had specified manually.

The optimiser is also useful for performing rapid reconnaissance of an area of the design space, followed by manual search to utilise know-how which is not so easily captured in a numerical performance measure. Sometimes these experiments reveal surprises about the product and its behaviour, as Figure 4 illustrates. Here the optimiser has output a sequence of six designs with steadily improving performance measures. Although the placement of the enclosure for design no. 6 is within constructable limits, the designer would prefer it further to the right and cautiously adjusts it in this direction (design moves towards the horizontal in this representation) to produce design no. 7. This degrades the performance a little so the optimiser is run, hoping to improve it. In fact, the result of this second optimisation, design no. 8, has far better performance than the previous best and, by way of a bonus, a much more acceptable positioning of the enclosure within the furnace walls.

A further investigation reveals that the slightly different roof

geometry of these two variants 6 and 8 is not the significant factor, but rather it is the additional downward movement of the enclosure that is compensating for its shift to the right. Restarting the optimiser at design no. 7 has located a new minimum, quite distinct from design no. 6, thereby showing an effect that could easily have been missed without the design steering and capture facilities.



**Figure 4** Combined automatic and manual search leads to a final design with excellent performance and desirable geometrical properties. Note that in this representation, shifting the enclosure to the right causes the design to move towards the horizontal, whilst shifting it downwards moves the point to the left.

## 5 CONCLUSIONS AND FURTHER WORK

We have described a new paradigm for the investigation of concept designs, called design steering, which we support by a novel representation of the abstract design space spanning the characteristic parameters of the product. Specification and review of designs is done from within this space, thereby reinforcing the sense of navigation of the space in order to search for optimal designs.

To relieve the designer of the chore of specifying large numbers of designs during a search, we have implemented an automatic optimisation facility. However, recognising that the best such design found in this way might not always be the most manufacturable, we have included a combined manual and automatic search capability. These features, along with the rapid delivery of design simulation results, promote experimentation with the design and a commensurate improvement in its understanding. In particular, the combination of a local optimiser and an informed manual choice of starting point appears to be an efficient method of finding good designs.

In the future we intend to apply the system to the heart valve design problem mentioned earlier. This problem has a larger number of parameters to be varied simultaneously, requiring multiples of the basic six-dimensional space to be displayed. We anticipate using a tiling of displays, such that the designer can choose to work on any one, but in the context provided by the remainder. Improvements in the labelling and initial setting up of the representation will also be needed as problem size increases. Finally, by extending the recording capabilities already present in the system, we intend to develop a collaborative version for use by a design team, possibly working at geographically separated locations.

## Acknowledgements

This work was carried out in part under the DIVA project at the University of Leeds (Design with Integrated Visualization and Analysis), jointly funded by the UK Engineering and Physical Sciences Research Council and the University of Leeds Academic Development Fund.

Grateful thanks are due to Ahmed Fazardi and Jerald Xu for the radiation gas furnace application. Also, to our collaborators Colin Shipway at BG plc (formerly British Gas plc) and Jeremy Walton at NAG Ltd we convey our thanks for the support and encouragement given by them and their colleagues. Finally we thank the anonymous reviewers for a number of helpful suggestions for improvement.

## References

1. M. Benedikt. Cyberspace: First Steps. In M. Benedikt, editor, *Cyberspace: Some Proposals*, pages 119-224, 1991. MIT Press.
2. S. Benford, J. Bowers, L. E. Fahlen, J. Mariani and T. Rodden. Supporting Cooperative Work in Virtual Environments. *The Computer Journal*, **37**, pages 653-668, 1994.
3. K. W. Brodlie, A. Poon, H. Wright, L. Brankin, G. Banecki, and A. Gay. GRASPARC: A Problem Solving Environment Integrating Computation and Visualization. In G.M. Nielson and D. Bergeron, editors, *Proceedings of IEEE Visualization 1993 Conference*, pages 102-109, Los Alamitos, CA, 1993. IEEE Computer Society Press.
4. T. David and C-H. Hsu. The Integrated Design of Mechanical Bi-leaflet Prosthetic Heart Valves. *Medical Engineering and Physics*, **18**, 6, pages 452-462, 1996.
5. A. Goel, C. Baker, C. A. Shaffer, B. Grossman, R. T. Haftka, W. H. Mason and L. T. Watson. VisCraft: A Multidimensional Visualization Tool for Aircraft Configuration Design. In D. Ebert, M. Gross and B. Hamann, editors, *Proceedings of IEEE Visualization 1999 Conference*, pages 383-387, San Francisco, CA, 1999. IEEE Computer Society Press.
6. Y. Jean, T. Kindler, W. Ribarsky, W. Gu, G. Eisenhauer, K. Schwan and F. Alyea. An Integrated Approach for Steering, Visualization and Analysis of Atmospheric Simulations. In G.M. Nielson and D. Silver, editors, *Proceedings of IEEE Visualization 1995 Conference*, pages 383-387, Atlanta, GA, 1995. IEEE Computer Society Press.
7. C. Johnson, S. G. Parker, C. Hansen, G. L. Kindlmann, Y. Livnat. Interactive Simulation and Visualization, *IEEE Computer*, **32**, 12, pages 59-65, 1999.
8. G. Kerlick and E. Kirby. Towards Interactive Steering, Visualization and Animation of Unsteady Finite Element Simulations. In G.M. Nielson and D. Bergeron, editors, *Proceedings of IEEE Visualization 1993 Conference*, pages 374-377, Los Alamitos, CA, 1993. IEEE Computer Society Press.
9. R. Marshall, J. Kempf, S. Dyer and C. Yen. Visualization Methods and Simulation Steering for a 3D Turbulence Model for Lake Erie, *SIGGRAPH Computer Graphics*, **24**, 2, pages 89-97, 1990.
10. J. D. Mulder and J. J. van Wijk. 3D Computational Steering with Parametrized Geometric Objects. In G.M. Nielson and D. Silver, editors, *Proceedings of IEEE Visualization 1995 Conference*, pages 304-311, Atlanta, GA, 1995. IEEE Computer Society Press.
11. J. A. Nelder and R. Mead. A Simplex Method for Function Minimization. *The Computer Journal*, **7**, pages 308-313, 1965.
12. A. J. Nowak. Solving Coupled Problems Involving Conduction, Convection and Thermal Radiation in Boundary Element Methods. In L. C. Wrobel and C. A. Brebbia, editors, *Heat Transfer*, 1992. Elsevier.
13. W. H. Press, S. A. Teukolsky, W. T. Vetterling and B. P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*, 2nd Edition, 1993. Cambridge University Press.
14. C. A. Shaffer, D. L. Knill and L. T. Watson. Visualization for Multiparameter Aircraft Designs. In D. Ebert, H. Hagen and H. Rushmeier, editors, *Proceedings of IEEE Visualization 1998 Conference*, pages 491-494, Research Triangle Park, NC, 1998. IEEE Computer Society Press.
15. L. Tweedie, R. Spence, H. Dawkes and H. Su. Externalising Abstract Mathematical Models. In *Proceedings of CHI '96*, pages 406-412, 1996. ACM Press.
16. H. Wright, K. W. Brodlie and M. J. Brown. The Dataflow Visualization Pipeline as a Problem Solving Environment. In M. Göbel, J. David, P. Slavik and J.J. van Wijk, editors, *Virtual Environments and Scientific Visualization '96*, pages 267-276, 1996. Springer-Verlag.
17. H. Wright, G.A. Stead and K. W. Brodlie. Interactive Exploration of Chemical Reaction Mechanisms using Novel Visualization and Integration Techniques. In M. Göbel and H. Müller and B. Urban, editors, *Visualization in Scientific Computing*, pages 166-173, 1995. Springer-Verlag.