

The Facilitation of Insight for Analog Design

Robert Spence, *Life Fellow, IEEE*

Abstract—The design of an electronic circuit or system—and, indeed, of many other artifacts and schemes—depends heavily upon the expertise, knowledge, and innate skill of the human designer. Unfortunately, many computer-based tools fail to fully exploit such human qualities. Thus, while a massive amount of research has been carried out in the essentially technical aspects of circuit design, very little effort has been directed to the support of essentially human aspects of design, such as the development of insight through exploration. This paper describes a new concept, embodied for illustration in a tool called the Influence Explorer, which facilitates human visualization of the relations existing between the parameters and performances of a circuit, thereby enhancing the designer's acquisition of insight into those relations and, as a consequence, the quality of a final design. The new tool is not intended to replace existing computer-aided design tools, but rather to complement them by exploiting a human activity largely unsupported at present.

Index Terms—Analog design, Influence Explorer, insight, prosecution matrix, qualitative exploration, visualization.

I. INTRODUCTION

A HUGE and rapidly growing volume of literature is devoted to the essentially technical aspects of circuit behavior. Theorems which identify limits to circuit performance and algorithms for the automated improvement of a design are just two examples.

By contrast, the typical circuit designer with a necessarily incomplete understanding of circuit design in general, and in particular of the circuit currently being developed, can benefit significantly from any insight gained into the complex relations between circuit parameters and circuit performances. Such insight might constitute awareness, for example, of the existence of a tradeoff between two critical performances or the severe effect, upon manufacturing yield, of high values of a particular parameter. This essentially human aspect of circuit design has been commented upon by many eminent circuit designers. Barry Gilbert, for example [1], remarks that

“The questions ‘What if ...?’ ‘Why not ...?’ and ‘How about ...?’ are often as important as ‘What?’, ‘Why?’, and ‘How?’”

II. EXPLORATION

The acquisition of insight and the search for possible circuit improvement is normally sought through exploration, by asking “What if?” questions of a simulator, usually by

Manuscript received October 28, 1997; revised December 3, 1998. This work was supported by Philips Research U.K. and by the U.K. Engineering and Physical Sciences Research Council.

The author is with Imperial College of Science, Technology, and Medicine, London SW7 2BT, U.K.

Publisher Item Identifier S 1057-7130(99)02353-8.

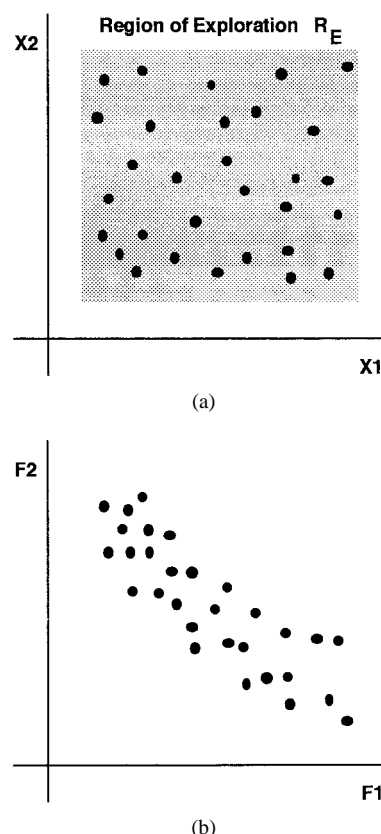


Fig. 1. The generation of many parameter-performance pairs by simulation.

changing the value of a parameter and observing the effect upon performance. Unfortunately, two drawbacks immediately present themselves. One is the delay between the variation of a component parameter and the display of its effect, a delay normally arising from the detailed nature of the model used by the simulator. A second is the inevitable result of the unavailability of a computable *inverse transformation* leading from desirable performance to suitable parameter values and tolerances. For this latter reason, the discovery of a tradeoff between two performances, for example, can be a long and tedious activity.

A principal aim of the tool described in this paper is to allow a circuit designer to visually explore the complex relations between parameters and performances *interactively*; with sufficiently rapid response, that insight can easily be acquired. Especially valuable, however, is the fact that it allows *performance* change to be the manually specified *cause*, and other performances as well as parameters to be the *effect*, thereby immensely facilitating the discovery of tradeoffs (between performances), as well as correlations. Following the acquisition of insight through such qualitative exploration, the




Fig. 2. Transformation from desired range discloses.




Fig. 3. Sweep of an F1 suitable parameter values a tradeoff with F2.

same tool then facilitates a smooth transition to quantitative design decisions.

III. PRECALCULATION

We shall illustrate the basis of the new design tool by means of a simple circuit described by two parameters ($X1$ and $X2$) and exhibiting two performances ($F1$ and $F2$) of interest to a customer. First, the designer identifies quite wide ranges [Fig. 1(a)] of the parameters $X1$ and $X2$, which together define a region of exploration (R_E) of parameter space within which a satisfactory circuit design is considered likely to be found. (Note that the boundaries of R_E are *not* related to parameter tolerances.) Next, this two-dimensional (2-D) region of exploration in parameter space is randomly sampled, typically identifying 200 or more designs within that region. The randomly selected designs are then simulated to identify, for each design, its location [Fig. 1(b)] in performance space.

Following this *precalculation* phase, the circuit designer now has available a database of designs sufficiently extensive to allow reasonably representative exploration of the relations between the two performances and the two parameters. For example, the designer may identify a rectangle in performance space (Fig. 2) defined by the customer's limits on the two performances and, immediately, note those designs in parameter space corresponding to performances within that rectangle. Alternatively, limits corresponding to a narrow range of $F1$ (Fig. 3) can be moved together over the entire range of $F1$ and the corresponding trend in $F2$ noted, in this case demonstrating a tradeoff between $F1$ and $F2$ and, perhaps, a correlation between $F1$ and a parameter. The advantage gained from precalculation is emphasized in Fig. 4; the interactive exploration loop typically used to acquire insight *now involves no simulation*, and can therefore be *responsive* in the sense that cause and effect are virtually simultaneous.

IV. PROBLEMS

Despite the simplicity of the concept just illustrated, two apparent drawbacks can immediately be identified. First, the number of circuit simulations is high, and would conventionally be considered prohibitive. Second, it is not clear how the concept illustrated in Fig. 1 could be extended to the realistic situation of *many* parameters and *many* performances; in this more general and realistic situation, the designer would be faced with understanding the nature of the distribution of a *cloud* of many designs distributed in two high-dimensional spaces, a task which could easily be as difficult as that of understanding the behavior of the original circuit.

V. CIRCUIT MODELS

The first perceived drawback is addressed by the availability of a class of models whose lack of complexity renders the simulation of hundreds of designs a feasible proposition. Macromodels [2], behavioral models [3], and other models of similar simplicity can be derived for a variety of circuits, and can thereby facilitate rapid simulation. In the design tool described below, we have in fact used response surface models [4] (see Appendix) which express each circuit performance (e.g., DC gain, bandwidth, slew-rate) as a relatively simple mathematical function of circuit parameters (e.g., resistance, capacitance). Nevertheless, certain insights might only be gained by—and therefore, fully justify the cost of—simulation at the transistor level.

VI. HISTOGRAM APPROACH TO HIGH DIMENSIONALITY

Extension of the concept illustrated in Fig. 1 to the case of many parameters and many performances is achieved by the choice of presentation shown in Fig. 5. This figure shows two sets of histograms, one for parameters and the other for performances. The data presented is derived from 160 simulations of the circuit of Fig. 6 using a response surface model derived by the process described in the Appendix. The five parameters are the aspect ratios (W/L) of five devices, and their values were selected randomly within the very wide exploratory ranges indicated in Table I.

Each design, therefore, is represented by a small rectangle within each histogram, as illustrated in Fig. 5 in response to an interrogatory click on one histogram. A fundamental advantage of the histogram presentation is that it provides a satisfactory solution to the problem of dimensionality; the existence of more parameters and performances simply requires the use of more histograms. Most crucially for design activity, use of a histogram set places equal perceptual-presentational weight on each parameter and performance, an essential property that would be difficult or impossible to achieve by three-dimensionality, three-dimensional (3-D) families [5], and other ways of increasing the dimensionality beyond two. An *apparent* drawback—that screen space will eventually limit the number of parameters and performances that can be handled—will later be seen to be a relatively minor problem.

Having shown in outline how the apparent disadvantages of the concept shown in Fig. 1 can be ameliorated, we now

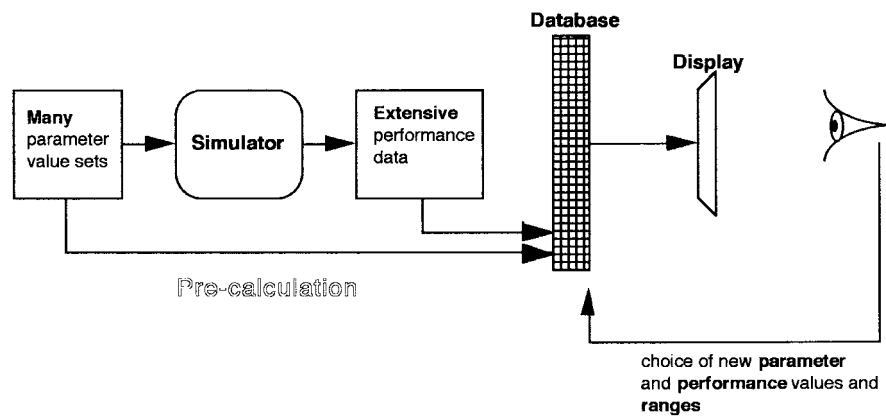


Fig. 4. Many potential designs are pre-calculated. The iterative and highly interactive loop facilitating “What if?” exploration and design now involves no simulation. It therefore allows the instantaneous display of cause and effect.




Fig. 5. A set of parameter histograms (right) and a set of performance histograms (left). Each design is represented once on each histogram.




Fig. 6. The circuit for which precalculation was performed, leading to the histograms.

TABLE I
THE REGION OF EXPLORATION IN PARAMETER SPACE SHOWN IN FIG. 5

parameter	Limits to Region of Exploration	
	min	max
b_1	5	15
b_3	5	15
b_5	3	7
b_6	100	200
b_7	5	100

describe the visualization tool, the Influence Explorer, that is thereby made feasible.

VII. INFLUENCE EXPLORER

The Influence Explorer [6] exhibits many features and offers many benefits. The features and anticipated benefits are discussed below, roughly in the order in which they might be encountered during the design process.

A. Interaction, Encoding, and Linking

As just discussed, histogram presentation addresses the issue of dimensionality. The other major problem, that of helping the designer gain insight into distributions in two spaces of high dimensionality, is addressed by two closely related features: one provides the designer with an interaction mechanism permitting the definition of limits to parameters and performances, while the other provides linking between histograms. Both are illustrated in Fig. 7. In this figure,




Fig. 7. A first indication of a possible tradeoff between DC gain and phase margin, and an inverse correlation with parameter b_7 .

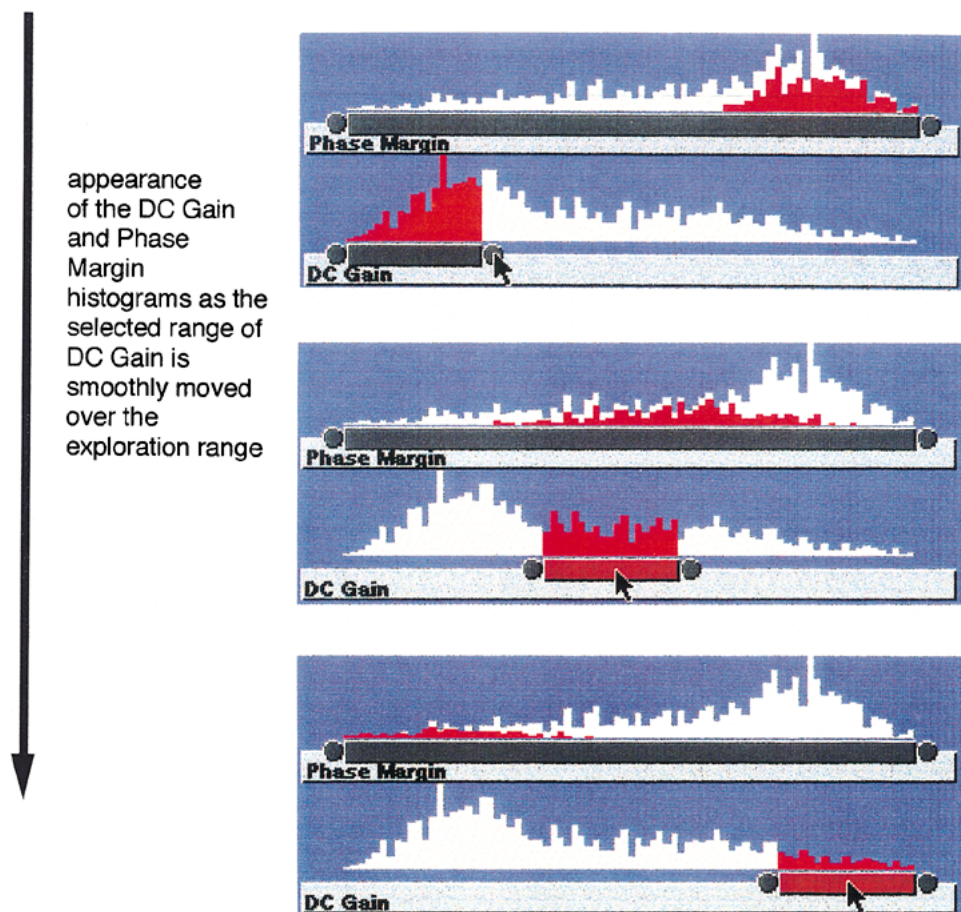


Fig. 8. Tradeoff discovery by exploration (view in sequence).

showing a subset of the histograms, the designer has adjusted the upper and lower limits of DC gain to define a range of DC gain of interest. This action identifies, in the precalculated database, all those designs (among the 160 samples) having a

gain between the two limits. The immediate effect is that the relevant part of the histogram is identified by color. But by far the most useful effect is that *these same designs are similarly color-coded on all other histograms*. The significance of this




Fig. 9. The saving of screen space by the encoding of summary information.




Fig. 10. The identification of limits to performances.

linking, which is easy to implement, is profound. Thus, as shown in Fig. 7 using selected performances and parameters, the designer might surmise that a tradeoff exists between DC gain and phase margin, and that there may well be a strong inverse correlation between gain and the parameter b_7 . Such a hypothesis can readily be checked by interactively moving the selected range of DC gain up and down the scale by simple movement of the bar between the limits and simultaneously observing the movement of the red parts on the other histograms. Fig. 8 illustrates this confirmation of the suspected tradeoff between DC gain and phase margin. The advantage of such a responsive exploration facility is difficult to appreciate from reading a textual description such as this one, but is immediately and strikingly apparent when actually used. Such visual immediacy allows a designer to quickly discover tradeoffs and correlations that would require extensive work using conventional tools, and generally gain insight into the complex relations between parameters and performances.

A question that immediately arises, and which is addressed in detail in Section VII-B, is that of *coverage*, especially as the number of parameters increases. It will later be shown how an automatically repeated precalculation can ensure an adequate sampling density as the parameter ranges are narrowed and necessarily fewer samples remain unless more are generated. Other important sampling issues are discussed in Section VIII.




Fig. 11. Reflection of performance limits into a parameter histogram.




Fig. 12. The distribution of pass and fail circuits.

B. Summary Encoding

The use of color to denote significant subsets within a histogram is one example of encoding. Another example, illustrated in Fig. 9, offers a useful approach to the problem of *too many* parameters and performances. On the assumption that designers will at any instant be focusing on a small number of parameters and a small number of performances, but nevertheless need to keep an overall eye on all variables, those that are not the current focus of attention can have the histograms removed, but the average value of the identified designs indicated by a yellow dot, as shown in Fig. 9; other properties, such as variability, could easily be encoded. The

advantage of such an encoding scheme is that any sudden change or trend exhibited by a variable not receiving focused attention would be immediately apparent and, if appropriate, the relevant histogram could be restored for more detailed examination. In this way approximately 30 parameters and 30 performances could be accommodated.

C. Qualitative Exploration

As described so far, the Influence Explorer is supporting qualitative exploration, hopefully leading to the acquisition of insight as opposed to the choice of specific parameter values leading to an acceptable design. However, a valuable aspect of the tool is that it facilitates a smooth transition, in both directions, between these two activities.

D. Customer's Requirements

The customer's preferred limits on performance are entered using the same mechanism as before. Thus, as illustrated in Fig. 10, upper and lower limits can be placed on all performances. If, for any reason, the designer does not wish to place such limits, they can be left at their extreme values.

E. Sensitivity Encoding

Fig. 10 also shows another valuable form of color coding. Any design that satisfies all the performance limits is coded red, as before. Any that fail only one limit are coded black, while dark grey and light grey indicate two and three failed limits, respectively. The advantage of such coding lies in the knowledge that if a limit is extended to include a previously black design, it will immediately turn red. In fact, two benefits are provided by this *sensitivity* display. First, some idea can be gained about the severity of the various limits; thus, (Fig. 10) in the illustrative example, the upper limit on phase margin is quite insensitive in the sense that extension of that limit will only include designs that have already failed other limits. Such information can, for example, facilitate negotiation with a customer in the event that a performance limit (which may have been quite informally specified) is unduly constraining design and might, perhaps, be relaxed. The other benefit arises when specified performance limits are so tight that no red components are shown; in this situation, black items give some idea of the extent to which limits have to be relaxed to include some satisfactory designs.

F. Parameter Design

The specification of performance limits, leading to the display of Fig. 10, also introduces color coding into the parameter histograms, as shown for a sample parameter in Fig. 11. The same color coding is used as for Fig. 10. Thus, selection of one of the *red designs* will ensure a circuit whose performances all lie within the customer's specifications. Such information is of limited value, of course, in view of the tolerances typically associated with parameters. Indeed, if the circuit design is to be mass-produced (and even if it is not), a more appropriate display of the parameter histograms may be, as shown in Fig. 12; here, any design that fails one or more

performance limits is encoded blue. This presentation now allows the designer to choose, within the very wide exploratory ranges of the parameters, a nominal value and tolerance (equivalently, upper and lower limits) for each parameter.

G. Design for Mass Production

The way in which the Influence Explorer is used from this point onwards will depend very much upon the needs of the designer but, to illustrate the power of the tool, we shall assume that a mass-produced discrete circuit is being designed and that the aim is to combine high manufacturing yield with low-cost components, the latter property associated inversely with the tolerance value. The manufacturing yield for the situation shown in Fig. 10, with all parameters randomly varying between extremely wide limits, will be extremely small. Remember, however, that these limits do not reflect manufacturing variability but rather the exploratory range suggested by the designer. What the designer may now do is select limits to parameter values in such a way as to maximize the estimated manufacturing yield. As limits to parameter values are changed, the effect is reflected in the red and blue color coding (Fig. 12), as well as in a displayed numerical value of the yield estimate. Thus, the distribution of red designs in the b_7 histogram of Fig. 12 suggests a low nominal value of that parameter. During the process of selecting parameter limits, the parameter histograms might appear as in Fig. 13.

In the illustration presented here, it has been implicitly assumed that the manufacturing distribution of the parameters is uniform, as was the case for qualitative exploration. In general, this will not be the case, and the precalculation must recognize appropriate parameter distributions when design for manufacturing yield is being carried out.

H. The Prosection Matrix

An alternative display can assist the designer in visualizing the location of a toleranced design within multidimensional parameter space. It is called the prosection matrix [7] and, for the illustrative example, is shown in Fig. 14. The color encoding will already be familiar from the earlier discussion. Basically, it shows (in yellow) how the parameter ranges, which together define the location of all manufactured designs in parameter space, relate to that (red) part of parameter space associated with acceptable circuits.

The way in which each rectangular display within the matrix is generated is shown in Fig. 15 for the simple case of three parameters. For the p_1/p_2 plot, the selected range of p_3 defines a section of multiparameter space. Those randomly generated designs which fall within that section are then projected (hence the term "prosection" [8]) onto the exploratory region of the p_1/p_2 plane. Thus, the top left component of Fig. 14 shows, projected onto the b_1/b_3 plane, all those designs lying within the tolerance ranges chosen for b_5 , b_6 , and b_7 , each design being color coded in the same way as before. Also shown in yellow in the top left component is the boundary of the tolerance region in the b_1/b_3 plane. What we are seeing, therefore, are images describing the relative locations, in pa-

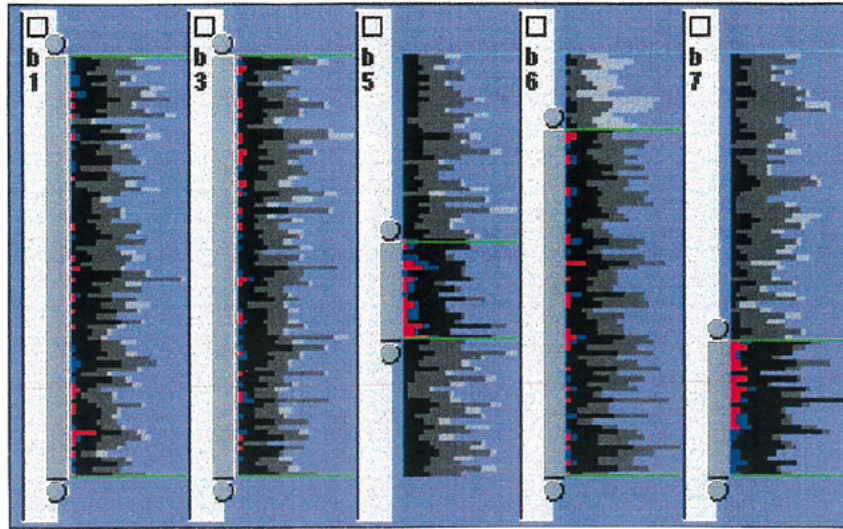


Fig. 13. Parameter limits are being selected to increase the manufacturing yield.




Fig. 14. The prosecution matrix for the circuit of Fig. 6.

parameter space, of the (yellow) tolerance region [9] describing the variability of manufactured circuits and the (red) region of acceptability [9] describing the location of acceptable circuit designs.

Two comments are particularly relevant. First, the underlying justification for the prosecution matrix (as for the earlier histogram sets) is that a difficult cognitive problem—that of choosing nominal values and tolerances leading to an economic design—has been transformed into a much easier *perceptual task*; in this case, that of making the yellow hypercube as large as possible while lying virtually entirely within the red region of acceptability. Second, each of the boundaries exposed by the color coding corresponds to a performance limit. Thus, an exploratory adjustment of a performance limit will cause the relevant boundary to move, thereby giving the circuit designer useful insight into how the various performance limits constrain the design [10].




Fig. 15. Construction of a prosecution matrix.

I. Refinement of Design

As design proceeds, and qualitative exploration is gradually replaced by quantitative decisions, the parameter limits are usually narrowed to correspond to realistic tolerances. In this event the number of precalculated designs lying within the tolerance region will be small and insufficient to provide the accuracy needed for yield prediction. It is a straightforward matter to arrange that, as parameter ranges are narrowed, a new *precalculation* is carried out to provide an adequate density of designs within the tolerance region. In many cases, such a calculation can be carried out in *real time* as design proceeds.

VIII. SAMPLING

Much of the potential advantage of the Influence Explorer and prosecution matrix rests on the manner in which parameter space is sampled in the precalculation phase. Sampling exerts considerable influence on computational cost, the information gained per simulation, the accuracy of the manufacturing yield estimate, and the appearance and consequent value of the prosecution matrix. Uniform Monte Carlo sampling was




Fig. 16. The creation of an RSM.

used earlier for simplicity of illustration, whereas substantial enhancement of the two visualization tools can accrue from a number of modifications to the sampling scheme.

Recent studies [11] illustrated in the context of circuit design have shown that to achieve maximum information per run (e.g., per circuit simulation), random sampling (such as the Monte Carlo approach) should be replaced by a patterned set of experimental points. In fact, one particular choice of a patterned set leads to significant enhancement of the Influence Explorer and projection matrix. Sampling based on Graeco–Latin squares [12] ensures uniform frequencies of points along parameter axes and has been found [13] to lead to a drastic reduction in the number of circuit simulations required and elimination of the problem of “unsampled gaps” in projections.

Different sampling distributions will almost certainly be pertinent to the two activities of qualitative exploration and quantitative design. During exploration, with all its implied uncertainty, it is not unreasonable to employ uniform sampling, as in the illustrations. However, during quantitative design, the sampling distribution must accurately reflect the manufacturing process if accurate yield estimates are to result. As was already explained in Section VII-I, the accuracy of a yield estimate, as well as coverage of the region of exploration, can be enhanced by carrying out, automatically, new “precalculations” as parameter ranges are narrowed by the designer.

It is also pertinent to point out that the visualization tools described in this brief can always exploit any prior knowledge of parameter correlations and other statistical information.

IX. EVALUATION AND GENERICITY

The Influence Explorer and projection matrix have been evaluated in two ways. First, a so-called *formative evaluation* [10], using a number of human subjects, was carried out by a psychologist and yielded useful data leading to improvement. Second, the new tool has been applied—in varying degrees of detail—to other fields of design, including structural design [14], financial design [10], and the design of electromagnetic

components. In all cases, the value of the Influence Explorer and projection matrix was commented upon favorably.

X. FURTHER RESEARCH

A. Flexibility

What the Influence Explorer and projection matrix have demonstrated is essentially an approach, based on precalculation and appropriate visualization techniques, to the acquisition of insight into the complex relations between parameters and performances. Within that approach, many different combinations of presentation, interaction, and linking are possible. For example, it would be entirely possible for each region of acceptability corresponding to a single performance to be displayed separately in the projection matrix and respond to changes in performance limits. Another facility, drawing upon the well-known parallel coordinate plot technique [15], could join designs on adjacent histograms by means of straight lines, thereby providing an alternative indication of correlations and tradeoffs and supporting the identification of convex and concave [16] segments of the region of acceptability.

B. Automated Design

The centers-of-gravity design centering algorithm [9] for the maximization of yield by suitable choice of nominal parameter values has been interfaced to the Influence Explorer. However, while the algorithm can be quite useful, its rapid execution, leading to little insight for the designer, emphasizes the need to provide an appropriate interaction [17] with the designer, perhaps one in which the algorithm is executed sufficiently slowly for the designer to gain insight into features of the design that are of interest.

C. Evaluation

The continuing evaluation of the Influence Explorer and projection matrix should include circuits whose performances are nonmonotonic with respect to parameters.

XI. CONCLUSION

A new design tool concept has been described which specifically addresses a circuit designer’s need to acquire insight through exploration. It is based on the precalculation of a sufficient number of designs to adequately characterize a region in design space within which a satisfactory design might be expected to lie, and the use of interactive linked histograms to support the designer’s exploration of the complex relations between parameters and performances. The tool is intended to complement, not to replace, conventional computer-aided design tools.

APPENDIX

RESPONSE SURFACE MODELS (RSM)

The aim of response surface modeling is to derive a simple mathematical model which adequately describes the behavior of an artifact over a selected range of the parameter values which together describe that artifact.

Briefly, the RSM methodology involves three stages, almost always repeated iteratively. First, parameter space is parsimoniously sampled (Fig. 16) according to some sampling scheme. For example, for a 2-parameter artifact, designs at the four vertices of the chosen exploration region may be identified. Next, the performance of the artifact at these sample points in parameter space is determined, either by simulation or by physical measurement. The third stage involves the proposal and testing of a mathematical relation designed to fit, as accurately as possible, the simulated or measured performances. Statistical expertise, as well as skill in modeling, is crucial to this third stage, which usually involves a number of iterations between the proposal of a model and its evaluation. Frequently the model is a polynomial expression, but the concept of generalized linear models [18] allows other forms of surface to be proposed and evaluated. It may well be decided, at this stage, that a denser sampling of parameter space is needed to derive a model of adequate accuracy, in which case (Fig. 16) the first stage is repeated. The interested reader is referred to the literature for methodologies [4], sampling schemes [19] and illustrative applications [4] of response surface models.

The form of a typical response surface model is illustrated by the model derived by the above process for the DC gain of the circuit of Fig. 6

$$\begin{aligned} \text{DC gain} = & 77.16 + 0.824b_1 - 2.428b_5 + 0.05b_6 - 0.4888b_7 \\ & - 0.0206b_1^2 + 0.1269b_5^2 - (9.8E - 5)b_6^2 \\ & + (2.48E - 3)b_7^2 \text{ dB.} \end{aligned}$$

Currently, the use of response surface models involves either delegation of the development of models to statisticians or for the designer to develop sufficient statistical expertise to carry out the modeling task. Severe disadvantages are associated with both these approaches, and it is for this reason that research is currently being undertaken to develop modeling tools that will allow the designer—the *problem holder*—to create satisfactory models.

ACKNOWLEDGMENT

The author wishes to acknowledge the kind and extensive support of the many colleagues he collaborated with and, as a result, benefitted immensely from. They include Dr. L. Tweedie, Dr. Z. Malik, Prof. J. Nelder, Dr. H. Su, A. Smith, H. Dawkes, R. Bhogal, and Prof. A. Brombacher, and additionally B. Gilbert and K. Chakhawata, whose interaction led to interesting insights.

REFERENCES

- [1] B. Gilbert, private communication, 1996.
- [2] G. R. Boyle, B. R. Cohn, D. O. Pedersen, and J. E. Solomon, "Macro-modeling of integrated circuit operational amplifiers," *IEEE J. Solid-State Circuits*, vol. 9, pp. 353–363, 1994.
- [3] T. Koskinen and P. Y. K. Cheung, "Hierarchical tolerance analysis using statistical behavioral models," *IEEE Trans. Computer-Aided Design*, vol. 15, pp. 506–516, 1996.
- [4] G. E. P. Box and N. R. Draper, *Empirical Model-Building and Response Surfaces*. New York: Wiley, 1987.
- [5] S. Feinder and C. Beshers, "Worlds within worlds: Metaphors for exploring n -dimensional virtual worlds," in *Proc. Assoc. Comput. Mach. 1990 Conf. User Interface Software Design*, 1990, pp. 76–83.
- [6] L. A. Tweedie, R. Spence, H. Dawkes, and H. Su, "The Influence Explorer," in *Proc. Assoc. Comput. Mach. Companion CHI'95*, p. 2.
- [7] L. A. Tweedie and R. Spence, "The projection matrix, a tool to support the interactive exploration of statistical models and data," *J. Comput. Stat.*, vol. 13, no. 1, pp. 65–76, Jan. 1998.
- [8] G. W. Furnas and A. Buja, "Prosection views: Dimensional inference through sections and projections," *J. Comput. Graphic Stat.*, vol. 3, no. 4, pp. 323–353, 1994.
- [9] R. Spence and R. S. Soin, *Tolerance Design of Electronic Circuits*. London, U.K.: Addison-Wesley, 1988.
- [10] L. A. Tweedie, "Exploiting interactivity in graphical problem solving: From visual cues to insight," Ph.D. dissertation, London Univ., U.K., 1997.
- [11] Z. Malik, H. Su, and J. Nelder, "Informative experimental design for electronic circuits," *Quality Reliability Eng. Int.*, vol. 14, pp. 177–186, 1998.
- [12] J. Dénes and A. D. Keedwell, *Latin Squares and Their Applications*. London, U.K.: English Universities Press, 1974.
- [13] A. J. Smith, J. A. Nelder, Z. Malik, L. Tweedie, and R. Spence, "Graeco-Latin square sampling for prosections," Dept. Elect. Electron. Eng., Imperial College, London, U.K., Info. Eng. Rep. 97/1, 1997.
- [14] H. Su, J. Nelder, P. Wolbert, and R. Spence, "Application of generalized linear models to the design improvement of an engineering artifact," *Quality Rel. Eng. Int.*, vol. 12, pp. 101–112, 1996.
- [15] A. Inselberg, "The plane with parallel coordinates," *Visual Computer*, vol. 1, pp. 69–97, 1985.
- [16] L. M. Vidigal and S. W. Director, "A design centering algorithm for nonconvex regions of acceptability," *IEEE Trans. Computer-Aided Design*, vol. 1, pp. 13–24, Jan. 1982.
- [17] L. Colgan, P. R. Rankin, and R. Spence, "Steering automated design," in *Proc. Conf. on AI in Design*. London, U.K.: Butterworth, 1991, pp. 211–230.
- [18] P. McCullagh and J. A. Nelder, *Generalized Linear Models*, 2nd ed. London, U.K.: Chapman & Hall, 1989.
- [19] W. G. Cochran and G. M. Cox, *Experimental Designs*, 2nd ed. New York: Wiley, 1968.

Robert Spence (S'55–M'59–SM'76–F'77–LF'97) has research interests in both engineering design and human-computer interaction. Following graduate work at Imperial College, London, he joined the Research Division of General Dynamics/Electronics in Rochester, NY. Since 1962, he has been on the academic staff of the Department of Electrical Engineering of Imperial College, where he is now Professor of Information Engineering and head of the department.

His research in engineering design is now directed to the use of response surface models, and the development of a tool which allows a designer with little or no knowledge of statistical techniques to create an acceptable response surface model. In the field of human-computer interaction, he was the co-architect of the MINNIE and CoCo circuit design tools and the co-inventor of the bifocal display, often referred to as the fisheye lens. More recently, he has been involved with the invention of novel information visualization tools. For his research in circuit design, he was awarded the Ph.D. (Eng) and D.Sc. (Eng) degrees and, for his contributions to human-computer interaction over 30 years, the Higher Doctorate of the Royal College of Art. He is a Fellow of the Royal Academy of Engineering.