Towards a Methodology employing Critical Parameters to deliver Performance Improvements in Interactive Systems

William Newman and Alex Taylor

Xerox Research Centre Europe
61 Regent Street
Cambridge CB2 1AB
UK
wnewman, ataylor@xrce.xerox.com

ABSTRACT In the design of interactive systems, overall performance issues tend to receive less attention than provision of functionality. As a result, systems may not always offer performance improvements to their users. In this paper we present some proposals for methods based on the use of application-specific *critical parameters*, i.e., performance parameters that measure how well the system serves its purpose. We argue that these parameters can provide a basis for dealing with performance issues in design. To construct this basis, there is a need to focus on identifying parameters through field studies, and on constructing models for use in making performance predictions. We provide a number of examples of critical parameters, and discuss problems we have encountered in identifying them and in model construction. We summarise the results of a study of designers working with critical parameters. A final section of the paper discusses the feasibility of introducing critical parameters to HCI design practice.

KEYWORDS performance metrics, critical parameters, methodology, interactive system design

1. INTRODUCTION

How can interactive systems be designed to deliver real performance improvements to their users? What new methods might help designers achieve this? The research we describe in this paper has been motivated by our perception that performance improvement is one of the most challenging issues facing Human Computer Interaction (HCI). This issue is certain to grow in importance as businesses make increasing demands for measurable returns on their investments in technology. The performance of interactive systems has been on the HCI agenda ever since Card, Moran and Newell began their pioneering work two decades ago (Card 1983). It has come to be a cornerstone in the argument for paying attention to HCI in system design.

In these circumstances, we would expect the delivery of performance improvements to the user to figure strongly in present-day HCI design methods, but this is not what we find.

To be presented at Interact '99 Conference, 30 August to 3 September 1999, Edinburgh, Scotland

Instead there is a lack of attention to performance improvement, noted by a number of writers. Landauer, for example, suggests that progress in most computer applications has been based on responding to market forces rather than on designing to meet productivity targets (Landauer 1995). In a commentary on HCI, Gibbs remarks on the lack of published papers offering evidence of improved performance in proposed new systems (Gibbs 1997). The relatively low priority given to performance issues in the packaged software industry has also been noted, for example in (Lammers 1986; Cusumano 1995).

Studies of designers at work likewise give the impression that performance issues rarely receive attention during design. For example, Herbsleb and Kuwana conducted a study of 38 design meetings and noted the types of questions that arose (Herbsleb 1993). There were so few performance-related questions that they did not merit treatment as a question type (J. D. Herbsleb, personal communication). In our own study of thirty designers' published accounts of their design projects, we found that performance criteria were rarely mentioned, and that opportunities to quantify performance, e.g., by setting targets or through user testing, were taken only rarely. When attention was paid to performance, the focus was rarely on the design as a whole, and almost always on a poorly performing component of the design: an error-prone selection technique, an easily forgotten command or a slow form of feedback, for example (Newman & Taylor 1997).

Many proposals have been made for improvements to standard HCI practice, but only a very few attempts have been made to help designers deal with issues of performance. Lim and Long's MUSE method, for example, includes performance specification within its scope, but its focus is primarily on a comprehensive method for the analysis and specification of functional requirements (Lim 1994). One major set of contributions has been the development of the GOMS cognitive modelling technique and its variants (Kieras 1997), which have proved their effectiveness but are still not widely used by practitioners. The cognitive walkthrough method has also demonstrated its usefulness in achieving designs supporting exploratory learning (Polson 1990); however, it is primarily a fault-detection method rather than a means of achieving learning-rate targets.

Despite this rather gloomy picture, we are encouraged by evidence that designers of interactive systems will pay attention to performance targets, if these are clearly stated and achievable. MacLean *et al*. demonstrated this in a study where they asked a pair of designers to design a "fast ATM" capable of speeding up cash withdrawals. Unlike the teams studied by Herbsleb and Kuwana, the two designers made frequent reference to performance criteria in the course of design. They appeared to have little difficulty in treating these criteria along with the rest of their concerns (MacLean 1991).

In this paper we describe an exploration of the use of *critical parameters* to define performance targets. This approach is attractive because it permits existing HCI design methods to be retained while introducing performance targets as additional design criteria, rather as MacLean *et al*. did in their study. We identify two main areas where research is needed, parameter identification and model building, and we summarise our progress in these two areas. We comment on the results of a study of designers working with critical parameters, and conclude with a discussion of some of the principal questions raised.

2. USE OF CRITICAL PARAMETERS IN DESIGN

Critical parameters are widely used in design as a basis for setting performance requirements. In general terms, they are the "established parameters by which designers measure whether an artefact or system serves its purpose, and compare one design with another" — a definition we have proposed in an earlier paper (Newman 1997). We went on to argue that the use of critical parameters would be beneficial to the design of interactive systems, because more attention would then be paid to performance enhancement and less to functional differentiation of products, which we claim is of less value to the user.

The use of critical parameters in interactive system design is attractive because it permits performance targets to be defined at the outset, before commitments have been made to solution strategies. Likewise when usability testing is done these targets can influence the choice of tasks to be performed, instead of *vice versa*. Altogether, use of critical parameters appears to improve the likelihood that overall performance will be taken into account during design.

We have taken the concept of critical parameters as a representation of system performance, and have begun to develop methods for identifying these parameters and for using them in design. In doing this, we have extended and strengthened our original view of critical parameters in several ways.

2.1 Critical Parameters: a Definition

We define a critical parameter as a metric for an aspect of the system's performance that is:

- **critical** to the success of the system in serving its purpose;
- **persistent** across successive systems for this purpose, and therefore valid as a basis for measuring performance improvements;
- **manipulable** by designers, who can design so as to achieve specific performance targets.

We have thus taken our original definition and added two properties that are essential to the purpose that we see these parameters serving. Performance parameters that are *persistent* not only allow measurement of improvements; they avoid the need to identify design criteria afresh every time a system is developed. Instead the same parameter can be applied as before, setting an appropriate new value as the target.

Parameters that are *manipulable* by designers can influence design decisions directly, during design iteration, rather than after a prototype has been built, tested and found wanting. Techniques such as GOMS can be used to simulate usage and make predictions of performance. Of course these techniques can be applied to any design target, such as the time taken to type a command or resize a window. If the target is critical to the design's overall success, however, performance predictions will enable the designer to manipulate the overall outcome.

2.2 An Example: Reviewer Assignment Time

In the organisation of academic conferences, a central responsibility of the programme chair is the distribution of papers to reviewers. The assignment process needs to be completed with the minimum of delay, but is inherently labour-intensive. It involves identifying, for each paper, a set of candidate reviewers who are expert in the necessary areas, and who have not already been assigned their full quota of papers to review. The chair then checks each candidate for possible conflicts of interest, and on this basis selects the requisite number of reviewers, making sure that each reviewer receives roughly the same number of papers. This process is repeated many times during the assignment process — if 400 papers are submitted to the conference and each requires five reviewers, 2,000 selections must be made. The process can therefore involve several days of a highly qualified academic chair's time, an expensive commodity.

An important parameter of this process is the *reviewer assignment time*: the time taken to assign a reviewer to a paper and to record the assignment. Other stages in the process are also time-critical, such as placing hardcopy papers in envelopes and mailing them to reviewers, but these do not depend solely on the programme chair. Assignment time is not only *critical* to completing the entire distribution process; it is *persistent* from one year's conference to the next, and it is *manipulable* through the design of assignment-handling procedures, whether paper- or computer-based. Hence reviewer assignment time can be regarded as a critical parameter in designing ways to support the assignment process. So too can error rate in recording assignments; see (Dix 1998) for an interesting study of reliability in paper reviewing processes. For the sake of this example we will focus on reviewer assignment time.

With this parameter in mind, various support systems can be envisaged and evaluated. An entirely paper-based system could be used, involving a printed master list of reviewers from which selections could be made on a basis of matching reviewers' interests to author-supplied keywords. Records of assignments could be kept manually for each paper and each reviewer. It takes only a cursory analysis to discover that, for a large conference, reviewer assignment time is likely to be very long if accurate matching is to be achieved. A semi-automated system, using a matching program to generate lists of candidate reviewers, could speed up the process, but reviewer assignment time is likely to remain high if the remaining steps are performed on paper.

A more efficient solution is therefore to implement a user interface of the kind shown in Figure 1, presenting the programme chair with a rank-ordered list of candidates generated by the matching program. The names in the list are checked for conflict, and the requisite number are selected and entered in a database, from which mailing labels can be generated later. This solution can be evaluated, for example with the aid of GOMS analysis, to arrive at an estimate of reviewer assignment time. It is feasible to work towards a specific target, such as 30 seconds to assign a set of 5 reviewers, by tuning the user interface and the underlying matching software.

3. THE METHODOLOGY OF CRITICAL PARAMETERS

Readers who find nothing very remarkable about the foregoing example of critical parameter usage are largely right. The route taken towards the design shown in Figure 1 involves no more than standard HCI principles and evaluation techniques, such as are found in Preece (1994). The significant difference is that a critical parameter is identified at the outset and is used to guide the iteration. Also, alternative designs can be compared for performance without building a working prototype, by using an analytical model. The two essential activities, identification of critical parameters and construction of models, are what we believe are missing from current practice. They form the basis of the new methods that we propose should be incorporated into HCI methodology. In the remainder of this section we explain our proposals in further detail. The paper then goes on to discuss practical aspects of parameter identification, model building and application of critical parameters.

Paper no. 1035
Telemedicine on the World Wide Web
Jane Smith
Health Sciences Dept, Tower Hamlets Univ., London

Candidate reviewers: **5 assigned**

[ ]	0.85	Peter Smith Tower Hamlets Consultants
[√]	0.81	Doreen Handscombe EE Dept, Driftwood College
[√]	0.72	Terry Johnson Tarragon Inc.
[ ]	0.71	Patricia Markham CS Dept., Tower Hamlets University
[√]	0.66	James Fallowfield MediOnLine plc
[√]	0.55	Charles Farmer Medical Research Horizons
[√]	0.54	Karl Haflinger Univ. Medical Centre, Kaiserslauten
[ ]	…	

ASSIGN

Figure 1. A hypothetical user interface to a reviewer selection system, geared towards reducing reviewer assignment time. Candidate reviewers are displayed in rank matching order. The programme chair selects those that do not appear to present a conflict of interest, and clicks on **ASSIGN** to store these assignments in a database.

3.1 Where Critical Parameters Are Found

Critical parameters are widely used in established engineering disciplines. They are fundamental, for example, to the ability of aeronautical engineers to design aircraft like the Boeing 777 that can be fully tested on the drawing board. In every discipline in which they are found, critical parameters are treated in much the same basic fashion. We have been able to learn how to introduce critical parameters to HCI by studying accounts of how they have arisen and how they are now used in other disciplines (Vincenti 1990). In this sense the methods we are proposing for use in HCI are tried and tested, but need to be adapted and integrated into general HCI design practice.

The biggest challenge in adapting critical parameters to HCI is to apply them in a wide range of computer-supported domains. Critical parameters are intrinsically *application specific*. A parameter identified as critical to one application, such as the support of the conference review process, is unlikely to be relevant to another, such as the support of radiologists. Our primary focus is on applying critical parameters in just such knowledge-intensive professional work settings as these, because we believe it is here that the most important and persistent performance challenges lie.

Our choice of professional work as a focus may seem unwise, because professionals have considerable autonomy and individual skill, allowing them to organise their work in a wide variety of ways. Nevertheless we believe there are grounds, both theoretical and empirical, for the existence of critical parameters in much knowledge-intensive professional work. Theoretical grounds lie in the role of *norms* in enabling people to plan their work in the context of deadlines, quality requirements, prior experience and division of labour. Empirical grounds can be found in the results of ethnographic studies of professional work, which provide evidence of people's performance goals and their strategies for achieving them (Orr 1996, Harper 1998). We have gained further encouragement from the results of our own studies of professionals at work, some of which we outline later in this paper.

3.2 The Elements of Critical Parameter Usage

Experience from the engineering world suggests there are three essential components to critical parameter use:

- **identification** of parameters in the chosen application domain;
- **model-building** to support predictions of performance;
- **application** of the parameters in design, as targets during both design iteration and prototype testing;

The methodology of critical parameters is fundamentally iterative, not just when they are used in design, but in all three of these respects. Thus a number of iterations may be required to identify a critical parameter successfully, as Vincenti has described in his account of research into aircraft flying qualities (Vincenti 1990). Iteration is also intrinsic to the process of developing successively better models to support performance predictions, a process familiar to most engineering researchers. We have begun to gain experience in iterative parameter identification and model-building, and to understand the close relationship between these two discovery processes.

4. PARAMETER IDENTIFICATION

In our 1997 paper, we suggested that the identification of critical parameters may prove time-consuming (Newman, 1997). Our recent experiences suggest, however, that parameter identification can sometimes be quite quick and straightforward: the parameters become obvious once the application domain has been identified. For example, we have recently begun to investigate the problem of designing tools to support foreign-language translators. Here a critical parameter, familiar to every professional translator, is the number of words translated per hour.

4.1 Model Building

Once a parameter has been found that appears to be critical and persistent, the next question is whether it is manipulable by designers. In other words, can a model be constructed by which performance can be predicted during design iteration? This involves gaining an adequate understanding of the structure of the activity to be supported, as it is performed in the 'real world.' We have applied relatively standard task analysis methods here, relying primarily on data gathered from field studies. The principal challenge is to develop a model that encompasses the work as a whole, rather than selected tasks. In simple cases, such as reviewer assignment, this is quite easy. In other cases we have been unable to develop an encompassing model, and have had to redefine the scope of the parameter itself. The next section provides examples of this.

To illustrate how models can be constructed iteratively to support critical-parameter based design in complex domains, we will use the well-known example of call handling by Toll and Assistance Operators (TAOs), as described by Gray *et al*. in connection with their Project Ernestine research (Gray 1993). Here the main critical parameters are well established as a set of times for handling roughly 20 different types of calls to the operator, including credit-card calls and collect calls. Various models of TAO call handling have been used in the design of workstations since the 1950s. More recently, Gray *et al*. developed a comprehensive CPM-GOMS model that they used with success to predict the performance of a new workstation design. The model offered a basis for extension, by those conversant with CPM-GOMS methods, in order to predict the performance of alternative workstation designs. For example, it allowed investigation of pre-recorded voice prompts and of voice recognition applied to caller utterances.

4.2 Setting the Scope of the Search

Parameters that seem obvious as measures of the performance of work do not always turn out to be critical, persistent and manipulable. A fresh search for parameters must then get under way; typically this involves ethnographic field study methods. We have conducted our studies by starting with the obvious parameter, and then either broadening or narrowing the domain of search. The two directions of study are best explained through examples. Our first example covers a successful attempt of ours to identify a critical parameter for the support of General Practice consultations; our second takes as its starting point the observed needs of researchers in libraries.

In the UK the work of the GP revolves around his or her consultations with patients. During the consultation the GP must obtain the history of the patient's problem, conduct any examinations necessary, make a diagnosis, decide on a course of treatment, issue a prescription and instruct the patient on how to manage the problem. All of this must be compressed into a period of 5 to 10 minutes. Since the 1980s there has been increased use by GPs of desktop computers, primarily to maintain patients' notes online and to issue prescriptions. The UK National Health Service has defined outline requirements for these computers, and subsidises the puchase by GPs of accredited computers.

An obvious critical parameter in supporting this work is overall consultation time, and any technology that could reduce this time while maintaining quality of service would be a clear design success. This is an intractable problem, however, not least because patients' perception of service quality is strongly and positively correlated with duration of the consultation. To find a more tractable performance measure we conducted field studies in two GP health centres. We videotaped over 60 consultations and interviewed four GPs, and then analysed the videotapes. After one or two false starts, we were able to identify a quite unexpected problem with tasks involving notes and prescriptions: if these tasks took longer than about 10 seconds to perform, the patient was likely to interrupt — to pick up the 'free turn' in the conversation created by the GP's silence. The interruption could lead to spending up to half a minute on a topic of little importance to GP or patient. On this basis we believe that a critical, persistent and manipulable performance parameter for GP support is the *proportion of accesses to pa-*

tient documentation that can be completed within 10 seconds. We believe that a system designed to perform well in these terms would assist the GP in managing consultations within the overall time constraints.

A major barrier to rapid identification of critical parameters lies, we believe, with correctly delineating the application domain. We learned this when we investigated one of the parameters we had previously thought to be critical, the time taken by a library researcher to copy a verbatim passage from a paper source document (Newman 1997). As O'Hara *et al*. have since pointed out, this kind of verbatim copying needs to be treated in the context of the use to which the verbatim material is put, e.g., entry in a notebook, storage in a database or pasting directly into a document (O'Hara 1998). The use of a new technology for copying, such as the overhead camera proposed by Newman, cannot be evaluated purely in terms of its copying speed. Instead the application domain needs to be widened, possibly to include the authoring activity for which the verbatim material is needed. We are currently engaged in investigating possible critical parameters in this wider domain of document authoring.

5. APPLICATION OF CRITICAL PARAMETERS DURING DESIGN

Research effort thus invested in identifying critical parameters and constructing models can pay off only if these parameters and models are useful in design. We have gained a limited amount of experience in the use of critical parameters, partly through applying them in our own design projects, partly through conducting experiments with design teams, and partly through the study of others' attempts to design with or without the help of critical parameters. We report here briefly on our experiments and retrospective studies.

5.1 A Design Experiment

We conducted a series of five design exercises, each involving a team of two designers. Three of the five teams included human factors experts, and all of them had experience in designing interactive systems. We gave each team the same design problem, in which the requirements included some performance targets defined in terms of critical parameters. To save them from having to design from scratch, we presented them with an outline description of an existing design. We asked them to critique this design in terms of the requirements, and then to proceed to redesign the system to overcome any identified shortcomings.

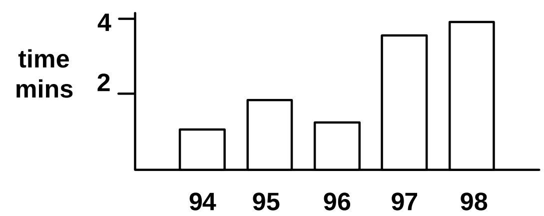

Figure 2. Variation in reviewer assignment time in five successive conferences. After the 1997 conference a set of web-based tools were used. Data supplied by programme chairs through personal communication.

We expected the inclusion of performance targets to assist the design teams in making real improvements. Instead, however, the teams appeared to gain little benefit from the parameters, in some cases ignoring them entirely, in other cases commenting on the need for a working prototype in order to test performance. Almost without exception, they devoted the design session to discussions of functionality, either of the existing or of the new design. One of the five teams did take up the challenge presented by the performance targets, and engaged in a rough GOMS-style analysis of the existing design. From this exercise we learned that it may not be realistic to expect designers to construct their own predictive models. We should perhaps have provided our design teams with a usable, tailorable analytical model to help them address performance targets.

In our studies of designers' past experience with critical parameters, we have identified a similar but more basic problem: designers need to be made aware of the parameters' existence, and constantly reminded of them. If they are not, they appear likely to resort to designing purely in terms of functional components. In a follow-up study of the review process, undertaken after identifying reviewer assignment time as the critical parameter, we gathered the data from which Figure 2 is constructed. It shows fluctuations in assignment time for a particular computer conference over a five-year period. A sharp increase occurred when new Web-based supporting technologies were introduced without taking assignment time into consideration.

A well documented example of the need to maintain awareness of critical parameters is found in Project Ernestine. The new workstation whose performance was modelled by Gray *et al.*, when compared with the workstation it was to replace, turned out to cause an increase in call-handling time. This was confirmed both by the CPM-GOMS model and by a four-month field trial. How could this have happened in the context of known critical parameters? Few details of the new workstation's design are available, but it appears that its performance was evaluated by means of usability tests in which users performed a series of timed tasks (Newman 1998). Each of these tasks formed part of a call, but no tests of complete call-handling were conducted. Thus the usability tests did not measure the workstation's performance in terms of the critical parameter. The design team was led to believe that performance had been improved, when it had not.

6. DISCUSSION

The concept of using critical parameters in HCI design raises a number of questions. Foremost among these are:

1. How widespread are critical parameters in the domains where interactive systems are used?
2. How persistent can we expect performance parameters to be in the face of rapid technological change?
3. Is performance, in terms of critical parameters, really manipulable by means of analytical models?
4. Can methods based on critical parameters be assimilated into HCI design practice?

The first question basically asks whether critical parameters perhaps exist in so few domains that the proposed new methods will rarely be used. We cannot offer hard evidence that critical parameters are ubiquitous; we can however report that, during just over a year spent searching for them, we have uncovered three sets of parameters and are hard on the heels of a fourth. We still rely to a large extent on arguments given earlier, concerning people's ability to plan their work and consistently achieve targets for completion times and reliability. If these people work to targets, it should be possible to identify related targets for the design of supporting systems.

At the same time, we recognise there are domains of work where people find targets difficult to set, and where we will probably find critical parameters elusive. Our own field of work, research, is one such domain. Indeed any work relying heavily on information search and retrieval is known to be unpredictable and hard to plan. This is one reason why the well-known critical parameters of information retrieval technology, precision and recall, have persisted in use for so long: it is very hard to identify metrics for the work in which these technologies are used. Pirolli and others have made significant progress in this direction in their work on models of retrieval cost structures (Pirolli 1995).

The second question concerns the persistence of critical parameters. Is there a real danger that parameters will turn out to be no longer critical after a system has been designed and introduced to the workplace? This could mean that effort spent on achieving performance targets is wasted. We accept that there are examples of performance measures becoming obsolete in the face of computerization, e.g., in the printing industry since the move to computer-based typesetting. However, these radical changes to the structure of work are usually followed by many years of incremental, measurable improvements.

We confess to being somewhat worried by Question 3, concerning the feasibility of building models that give accurate predictions of performance. We do not have sufficient experience to answer this question. We do believe, however, that the discovery of persistent critical parameters will encourage researchers to develop models and tools for performance prediction. As designers attempt to meet performance targets, and fail through a lack of analytical techniques, they will start to make demands on the research community.

As regards the last question, we have been concerned from the outset to propose methods that are easy for designers to incorporate in their current design practice. Designers will pay attention to performance issues if they are given clear and achievable targets, as MacLean *et al*. discovered. We have focused on finding ways to identify parameters that are sufficiently clear and achievable to deserve designers' attention. We have looked for ways to relieve the design team of the burden of identifying performance targets. We also hope we have reduced the need to factor usability testing into the design schedule.

7. CONCLUSION

We have devoted this paper primarily to discussing what critical parameters might mean for HCI design practitioners. We conclude by pointing out that there are implications for two other communities: researchers and users.

Use of critical parameters promises to introduce a new role for HCI researchers — parameter identification — and to raise the importance of a current role, the construction of models. These roles could provide opportunities to make even stronger contributions to design practice than at present. The contributions could take the form of enhanced models for predicting performance in terms of critical parameters; they might also include newly identified parameters. In this way the partnership between researchers, system designers and HCI practitioners could be further strengthened and enriched.

We hope that critical parameters also offer something new for users. At present, users can have little influence on the directions taken by software product suppliers, except when they switch allegiance en masse from one supplier to another. Critical parameters could provide users with a basis for demanding specific improvements in performance. This would have the beneficial effect of making suppliers more accountable for the performance of their systems, and ultimately might lead to competition on the basis of performance rather than just functionality.

Acknowledgements

We are grateful to our Xerox colleagues for providing constant feedback and encouragement as this work proceeded, and especially to Hervé Gallaire for initially engaging us in discussing the role of critical parameters. We would like also to thank the conference programme chairs who provided the data for Figure 2. Finally, we wish to acknowledge the many insights we have gained from Walter Vincenti, first from reading his remarkable book, and subsequently from a series of all-too-infrequent discussions.

References

Card S. K., Moran T. P. and Newell A. (1983) *The Psychology of Human Computer Interaction*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Cusumano M. A. and Selby R. W. (1995) *Microsoft Secrets*. Harper Collins.

Dix A., Ramduny D. and Wilkinson J. (1998) "Interaction in the large." *Interacting with Computers* Vol. 11 (1), pp. 9-32.

Gibbs W. W. (1997) "Taking Computers to Task" *Scientific American* July 1997.

Gray W. D., John B. E. and Atwood M. E. (1993). "Project Ernestine: Validating a GOMS Analysis for Predicting and Explaining Real-World Task Performance." *Human Computer Interaction* Vol 8, pp 237-309.

Harper R. H. R. (1998) *Inside the IMF: An Ethnography of Documents, Technology amd Organisational Action*. London: Academic Press.

Herbsleb J. D. and Kuwana E. (1993) "Preserving Knowledge in Design Projects: What Designers Need to Know." Proceedings of InterCHI '93 Human Factors in Computing Systems (April 24-29, Amsterdam, Netherlands) ACM/SIGCHI, N.Y., pp. 7-14.

Kieras D. E. (1997) "A Guide to GOMS Model Usability Evaluation using NGOMSL." In Helander M. G. and Landauer T. K. (Eds.) *Handbook of Human-Computer Interaction*, 2/e. Amsterdam: Elsevier.

Lammers S. M. (1986) *Programmers At Work*. Redmond, WA: Microsoft Press.

Landauer T. K. (1995) *The Trouble with Computers*. Cambridge MA: MIT Press.

Lim K. Y. and Long J. B. (1994) *The MUSE Method for Usability Engineering*. Cambridge: CUP.

MacLean A., Young R.M., Bellotti V.M.E., and Moran T.P. (1991) "Questions, Options and Criteria: Elements of Design Space Analysis." *HCI*, Vol. 6, nos. 3/4.

Newman W. M. (1997) "Better or Just Different? On the Benefits of Designing Interactive Systems in terms of Critical Parameters." *Proc. DIS '97, Designing Interactive Systems* (Amsterdam, 18-20 August), pp.239-245.

Newman W. M. and Taylor A. S. (1997) "Critical Parameters in Interactive System Design: Pointers from a Literature-based Study of Representations of Performance." Tech. Report EPC-1999-101, Xerox Research Centre Europe, Cambridge UK.

Newman W. M. (1998) "On Simulation, Measurement and Piecewise Usability Evaluation." *Human Computer Interaction* Vol 13 (3), pp. 316-323.

O'Hara K., Smith F., Newman W. M. and Sellen A. J. (1998) "Student Readers' Use of Library Documents: Implications for Library Technologies." *Proceedings of CHI 98 Human Factors in Computing Systems* (April 18-23, Los Angeles CA) ACM/SIGCHI, N.Y., pp. 233-240.

Orr J. (1996) *Talking about machines: an ethnography of a modern job*. Cornell Univ. Press.

Pirolli P. and Card S. K. (1995) "Information Foraging in Information Access Environments." *Proceedings of CHI '95 Human Factors in Computing Systems* (May 7-11, Denver, CO) ACM/SIGCHI, N.Y., pp. 51-58.

Polson P. G. and Lewis C. H. (1990) "Theory-based design for easily learned interfaces." *Human Computer Interaction*., Vol 5, pp 191-220.

Preece J., Rogers Y., Sharp H., Benyon D., Holland S. and Carey T. (1994) *Human Computer Interaction*. Wokingham: Addison-Wesley.

Vincenti W. G. (1990) What Engineers Know and How They Know It: Analytical Studies from Aeronautical History. Baltimore: Johns Hopkins Univ. Press