

Performance Targets, Models and Innovation in Interactive System Design

William M. Newman, Alex S. Taylor*, Christopher R. Dance and Stuart A. Taylor

Xerox Research Centre Europe

61 Regent Street

Cambridge, CB2 1AB, UK

+44 (0)1223 341 500

first.last@xrce.xerox.com

ABSTRACT

This paper presents an approach to designing interactive systems that enables critical performance parameters to be identified and models of performance to be constructed. The methods described are intended to enable designers to improve the performance of systems, and the provision of performance targets is expected to encourage innovation in design. An example is quoted in which digital camera technology was applied to the support of authors using paper source documents, to enable them to capture source text more rapidly and thus increase their productivity, measured in terms of words per hour. A model of the capture task was constructed, and was used to set a target time for capturing short text segments. This target was presented to a design team, who responded with an innovative interface incorporating auto-completion. A prototype auto-completion tool demonstrated that the performance target could be met.

Keywords

Camera-based scanning, critical parameters, auto-completion, innovation.

1 INTRODUCTION

Designers of interactive systems, like many others in the computer field, face the challenge of keeping pace with a rapidly advancing technology while delivering a stream of benefits to their clientele. In interactive system design the challenge is heightened by having user communities as clientele, and by having to pass benefits directly to the user. An understanding of how to deliver such benefits, and what methods to employ to ensure they are delivered, is fundamental to design success.

Over the last few years our own research has been concerned with developing methods by which designers of interactive systems can

deliver performance improvements to their users. We embarked on this research because we perceived that few such methods were available to designers. We found evidence of this in our studies of the HCI literature [9, 12]. We also found, elsewhere in the literature on design methods, insights into how methods for performance improvement might be developed [17,18].

Recently we have begun to focus on the potential role of application-specific *critical parameters* as a basis for measuring and predicting the performance of interactive systems [10]. These parameters provide designers with the means to assess how well a system serves (or will serve) the needs of its users, and to support the comparison of one design with another. We began simply by investigating whether these parameters might exist. We found from our initial studies that they did, and therefore began to consider how they might be used in design. In a recent paper we report on our investigations of critical parameters in a number of application domains [13].

A recurring issue in designing for performance is whether this promotes or discourages innovation and creativity. Specification of a performance target sets a constraint on the design and might therefore be expected to limit innovation. On the other hand, many examples can be quoted of highly innovative designs – the jet engine, the transistor, RISC technology, etc. – that have sprung from the need to overcome performance barriers. We would agree that setting generic performance targets for interactive systems could impede those searching for new applications for technology. But this is not what we are about; we want to understand how existing, identified applications can be better served by interactive systems. We claim that performance targets can indeed help designers achieve innovative solutions in specific application domains, and hope this paper helps justify this claim.

A second issue addressed in this paper is whether designers of interactive systems can readily work with critical performance targets. The results of recent research have been encouraging. We have conducted an experiment in which we first identified a critical parameter and used it to establish a performance target. We then passed this target to a design team, who were able to devise a solution that met the target. Furthermore the approach they took was relatively innovative: the challenge of meeting the specified target appeared to push the design team into considering fresh design options. In the final discussion section of this paper

* Current address: Digital World Research Centre, School of Human Sciences, University of Surrey, Guildford, Surrey GU2 5XH, UK; email: Alexander.Taylor@surrey.ac.uk.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

DIS '00, Brooklyn, New York.

Copyright 2000 ACM 1-58113-219-0/00/0008...\$5.00.

we consider some wider implications for the design of interactive systems.

2 CRITICAL PARAMETERS: OUR RESEARCH STRATEGY

Our research strategy hinges on demonstrating that critical performance parameters can be found in work activities. These parameters provide a basis for modelling the impact of interactive systems on the work, and for designing systems that support the work better. As explained in [13], we define a critical parameter as a metric for an aspect of the system's performance that exhibits three properties:

- It is **critical** to the success of the system in serving its purpose. The parameter therefore offers a basis for judging which of several solutions is superior. A well-documented example of a critical parameter is telephone operators' call-handling time, used by Gray *et al.* in their "Project Ernestine" study of two different operator workstations [3].
- It is **persistent** across successive systems designed for a particular purpose. It therefore avoids 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. Persistent parameters provide a basis for comparing successive designs and assessing cost/benefit. In the Project Ernestine example, call-handling time had been in use since the 1960s as a critical parameter in designing telephone operator workstations [6].
- It is **manipulable** by designers, allowing them to predict performance and thus make informed design tradeoffs that lead to meeting targets. Designers of telephone operator workstations, for example, aim to make changes to the user interface that reduce call-handling time; one outcome of Project Ernestine was to provide them with an improved model of call-handling that provided more accurate performance predictions. In situations where there is no such means of predicting performance, the designer cannot manipulate the parameter, and the result may disappoint.

We have identified critical parameters in a number of work settings. One of these, mentioned in [13], is reviewer assignment time in the conference paper review process. Several factors have motivated us to focus on work activities, including our own corporation's interest in developing technologies for the workplace. There are undoubtedly critical parameters in leisure activities too. Alm *et al.*, in their study of conversational prostheses, identified a critical parameter (words of speech per minute) that applies in any conversational setting [1].

2.1 Modelling performance

When we first embarked on this research, our interest was mainly to identify the parameters critical to design success. Now the primary focus of our research is increasingly on developing improved methods for *modelling* interactive system performance, with critical parameters as a basis for measurement. In this way we can equip designers with tools for making predictions of their designs' performance. In other words, we can ensure that the critical parameters we identify are manipulable.

The need for models to accompany parameters is illustrated in Figure 1. Parameters on their own enable designers and evaluators to focus on critical aspects of performance, but only when they have a working prototype available for testing (Figure 1a). By this time it is often too late to deal with a shortfall in performance. It may not even be clear how to deal with it. Armed with a model of the system's support for the work process, however, the designer can conduct repeated analyses of the system's performance whilst the system is still in an early stage of design, and can identify sources of shortfalls (Figure 1b). We believe that only a combination of metrics and models will overcome weaknesses in current system design methods.

We claim that knowledge of critical parameters can simplify the task of constructing predictive models. Rather than attempt to model all aspects of work performance, which would be out of the question, model-building can be focused on predicting just those aspects of performance that are critical. Our approach is therefore to look for a small number of critical parameters that provide measures of how well the work is performed overall. We construct models that enable designers of supporting technology to make predictions of performance, measured in terms of these critical parameters. The same parameters can be used when evaluating the performance of the finished system. In this way, designers can work towards meeting specific performance targets and can later test whether these targets have been met.

3 APPLICATION-SPECIFIC STUDIES

We have conducted a number of studies with a view to identifying critical parameters in the workplace. Our primary interest has been in supporting knowledge work. This is an attractive arena because the benefits to knowledge workers from interactive systems have been hard to measure in the past, and are likely to be increasingly important in the future. Our studies have covered several types of such work: primary health care, use of source documents in libraries, various kinds of authoring work and, most recently, information seeking in support of document creation. Our research methods vary from study to study, but typically involve two main study phases:

1. **Ethnographic fieldwork** directed towards identifying the work's overall structure and its critical parameters;
2. **Detailed data gathering and analysis** leading to modelling the effect of technology on overall performance.




Figure 1. Measuring and modelling performance: (a) using critical parameters CP_1 and CP_2 to measure a system's support for a work process; (b) using a model to make predictions of performance of the work process, in terms of the same critical parameters CP_1 and CP_2 . With the model, the designer can make performance predictions without the need for a working prototype.

These two phases draw heavily on conventional methods of field study and task modelling. In the first phase, however, we maintain a particular focus on gathering data that will help us to identify the structure of the work, gather examples that we can use as scenarios or benchmarks, and identify overall critical parameters. In our study of primary health care, for example, we videotaped about 80 consultations and followed them up with interviews. Our analysis enabled us to identify the principal stages of consultation (greeting, history taking, examination, diagnosis, treatment discussion, conclusion), consistent with others' findings [2, 5]. It also shed light on the ways patient records play into this sequence, an issue largely overlooked by previous studies. We were able to identify a critical parameter – time lost to disruptions in the consultation sequence – and to explore the influence of patient records systems on this parameter (Figure 2).

In our second phase of study we rely mainly on task analysis to discover elements of the work that influence overall performance. We have also found Conversation Analysis (CA) effective here. In the study of primary health care, for example, we were able through CA to perceive the potentially disruptive effect of physicians' accesses to patient records. If this access took more than about 10 seconds (as it often did when computer records were accessed) the patient was likely to break the resulting silence with a new and possibly time-wasting topic of conversation. We thus identified a property of the patient records system – percentage of accesses lasting longer than 10 seconds – which appears to influence the disruption-time parameter.

These studies are of necessity *application specific*. Knowledge work does not follow a single universal structure, nor can it be measured by a single set of critical performance parameters. Within a particular context, however, knowledge work often follows consistent patterns in both structure and performance. These patterns persist across time and can therefore guide the design of many generations of solution. In primary health care, for example, there is a persistent pattern to the consultation sequence (see above), and to the allocation of time to the consultation (7 to 10 minutes). This persistence of performance measures and processes could be maintained partly through consistency in the training each profession receives, and partly by requirements for work sharing. For example, in financial work an analyst comes to know that the data supplied by a colleague will be calculated in a particular way and will be supplied by a certain deadline. By focusing our studies on carefully delineated application domains we can discover these work structures and the parameters that go with them.

4 AN EXAMPLE: DESIGNING FOR RAPID SOURCE TEXT CAPTURE FROM PAPER

To illustrate the use of critical parameters we describe a project we have undertaken recently involving, as its application domain, the creation of documents using multiple sources. This domain had previously been under investigation for some time in our laboratory [14, 15]. We built on our earlier studies, applying critical parameter methods in order to design a supporting technology based on camera-based scanning.




Figure 2. Modelling the primary-care consultation process, and the effect of records system usage on time lost to disruptions.

4.1 Use of source documents while writing

We are concerned here with a work context in which an author is preparing a document, not entirely out of his or her own head, but instead relying partly on drawing text from existing documents. In the workplace this is an extremely common approach to document preparation. We have frequently observed authors working with a number of source documents, from which they extract segments of text, columns of figures, diagrams and images [15, 4].

Authors are especially prone to refer, while preparing an electronic document, to the *paper* form of a source document – a table of data, a report, a book, a photocopy and so forth. O'Hara and others [14] investigated this and other configurations of source and target documents, and found authors had a marked preference for working with source documents in paper form, rather than on-line, even though this might involve extra re-keying. In contrast, they preferred to create documents on-line rather than on paper.

Despite authors' almost universal preference for using paper sources during on-line document preparation, little attention has been paid to developing technologies to assist such authors. The technologies most widely recommended for text capture – hand-held and flatbed scanners – are cumbersome and slow to use. Their cost has fallen dramatically in recent years, yet there is little evidence that this is persuading authors to turn to them as a means of capturing information from sources while writing. Re-keying of paper-based source text is still universal.

4.2 Digital cameras as source capture devices

One device that has attracted our interest as a means of text capture is the digital camera. It lacks the flatbed scanner's capacity to capture large, high-resolution images of documents, but this disadvantage is rapidly receding as large sensor arrays become available at low cost. The main attractions of the digital camera are the speed with which it can capture images, and the convenience of a device that can be mounted over the desk where it can scan documents as they lie face-up.

A system developed at our research centre, *CamWorks*, uses a digital video camera for capturing text and images from paper documents [11]. The camera is mounted vertically over the desk beside the workstation, and a live video window displays the camera image to the user, enabling the document to be positioned under it (see Figure 3). The user can then select a portion of the document, using similar selection methods to those of a word processor, and can drag and drop the selected portion (text or image) into an electronic document.

Prior to our critical parameter study, usability tests had been conducted to compare text capture times for *CamWorks* with times



Figure 3: The CamWorks system in use. A digital video camera captures images of a document on the desk, from which the user can select text and image segments for copying into an electronic document.

for flatbed scanning and for retying the text. These are shown in Figure 4. They showed a reduction by approximately 40 seconds in capture time when CamWorks was used in place of the flatbed scanner. Re-keying time was of course dependent on the amount typed and on typing speed (here 35 words per minute). The results were encouraging, but could not confirm whether CamWorks would be effective for source usage while writing. We therefore undertook a series of studies, which we describe here as Studies A, B and C. They enabled us to identify critical parameters, construct performance models and thus design a new system whose performance we predicted (and later showed empirically) would make it preferable to other means of source text capture.

4.3 Study A: Establishing critical parameters for the authoring application

The main question that prompted the next phase of research, involving critical parameters, was this: could we develop a camera-based technology offering a viable alternative to re-keying source text from paper? We had convincing evidence that a digital camera could out-perform a flatbed scanner. We could not make any predictions, however, of the camera's ability to save the author from re-keying source text. The camera might well be rejected by users if they perceived that it supported them less well than the keyboard.

Our first step was to try to identify a critical parameter for measuring the technology's overall level of support to the author. We conducted a field study (Study A) involving ten professionals whose work involved large amounts of writing from sources [16]. They included consultants, lawyers, educationalists and academic researchers. Each of the participants was videotaped and observed performing a 'real-world' writing task in their everyday work setting. We learned from this study how the participants organised their source information, how they worked between multiple sources, and what steps they went through to construct

their documents. Importantly, we also learned how they gauged their productivity or performance. A critical parameter that emerged from this study was *words per hour drafted*. A number of participants confirmed that they used this measure of writing speed to estimate their performance.

4.4 Study B: Developing a model of authoring

Our next step was to develop a model of authoring that would explain and predict the technology's contribution to achieving words per hour rates. For this step we conducted a controlled study, Study B, involving six university students enrolled in a range of social science and humanities programmes. Using a simple experimental design, we asked each of the participants to write 500 words on a specific subject (human cloning), using a number of related sources which we provided on paper. We videotaped each session, and reconstructed the sequences by which each participant constructed their articles. In building this model we drew on previous studies of typing, in particular [7].

We found that most of the text copied from the sources was in the form of single words or short phrases, many of them terms, names and dates incorporated into paraphrases of the source material. With only six participants it was not possible to gain a reliable model of the distribution of phrase lengths, but there was a clear bias towards short phrases, nearly 60 percent of the phrases copied being of 4 words or less (see Figure 5).

4.5 Modelling text capture

The result of this second study was a partial model of how text capture contributes to writing speed. It provided a rough distribution of the word-lengths of captured text items, as shown in Figure 5. It also provided estimates of mean capture times for text segments of a given length (see Table 1). From the recorded times we calculated an approximate re-typing speed of 29 wpm. Taken together, these figures could have allowed estimation of the net gain to the user, across a writing task with similarly distributed captures, if capture times were altered. However, the small sam-




Figure 4. Average times to complete copying tasks using CamWorks, a flatbed scanner, and retying at 35 wpm using the PC keyboard (from [10]).




Figure 5. Distribution of average (over all authors) number of quotes against length of quotes measured in words.

ple size on which we based the measured distribution meant that we could not rely on this model in deriving critical parameters and targets.

We therefore adopted a simpler model, in which we relied only on reducing capture times across as wide as possible a range of text segment lengths. Based on experience with CamWorks we expected a significant reduction could be achieved in times for long segments, but short segment times might be hard to reduce. We adopted *short-segment capture time* as our critical parameter, and set the design team the task of matching the 8.2-second re-keying time for 4-word segments, and reducing times for segments longer than this.

4.6 The design of a camera-based text capture tool

We communicated the results of our study to the CamWorks design team. We explained the need to match the 8.2-second capture time for 4-word segments, as well as reducing times for the longer segments that authors captured less often. We gave them the data on observed capture times (Table 1). We reminded them that CamWorks, during evaluation, had exhibited a minimum time of around 25 seconds for very short captures; this would be very hard to reduce to 8 seconds.

The design team adopted the idea of using *camera-based auto-completion* to support faster captures. This was an idea that had arisen during our experiments with new techniques for interacting with video images of documents, but in the absence of motivating requirements it had not been tried or tested. With auto-completion, users could specify the text they wished to copy by typing the first few characters of the text, rather than by selecting it with the mouse. The user interface of the word processor could be modified to show, as the user typed, candidate words found in a document placed under the digital camera (see Figure 6). The user could accept the word using a special key, and could accept subsequent words by pressing the key again for each word. Alternatively the user could step to the next candidate word, or could simply ignore the offered completion and keep on typing.




Figure 6. (a) Portion of live video window displaying image of source. (b) Portion of electronic target document showing auto-completion from the source.

4.7 Modelling the design's performance

We made predictions of capture times using Kieras's Natural GOMS Language [8]. According to this model, if image processing and text recognition times could be reduced to zero, users would take approximately 3.6 seconds to copy one word using the auto-completion system. Each additional word would take a further 1.0 seconds to copy (Table 1). The auto-completion tool should thus be able to out-perform the speed of users re-keying 3-word segments at 29 wpm. It should offer an increasing advantage when the text segments were longer. We could expect users therefore to prefer the auto-completion tool for most if not all of their source captures. It seemed likely they would accept it, at least on performance grounds.

These predictions relied on reducing to zero the time taken to process the camera image and recognise the text, and this was patently impossible. In this initial design exercise, however, we were concerned less with reducing these computation times, and more with evaluating both the new design strategy and the whole approach of designing to meet targets designed in terms of critical parameters. We therefore chose to simulate, in the prototype system, the effect of 'instantaneous' image processing and recognition. We comment on this decision later, in the Discussion section.

Table 1: Actual re-keying times and predicted word-segment capture times using the auto-completion tool for a 29-wpm typist.

words	1	2	3	4	5	6	7	8	9	10
re-key, secs	2.0	4.1	6.1	8.2	10.2	12.2	14.3	16.3	18.4	20.4
prediction, secs	3.6	4.6	5.6	6.6	7.6	8.6	9.6	10.6	11.6	12.6

A prototype tool was built, driven by an auto-completion user interface embedded in Microsoft Word. Near-zero processing time was achieved by pre-processing all of the source documents, and printing on each one a unique barcode that enabled the auto-completion tool to retrieve and display the pre-processed contents

as soon as the document was placed under the user's camera. The prototype was subjected to a short usability test in which some user interface problems were identified; these were then resolved.

4.8 Study C: Validating the performance estimates

To validate our model and the auto-completion tool itself, the system was used in a final authoring experiment (Study C). Six students took part in this experiment, all drawn from the same group as before but none of them participants in Study B, and all having typing speeds of approximately 29 wpm. Each was asked to complete an authoring task identical to the first, but this time using the auto-completion tool. Prior to the experiment they were trained to use the tool until they reached a specified level of competence and appeared comfortable with its use. As before, the sessions were videotaped and the sequences of sentence construction were extracted.

Table 2 summarises the quantitative results of Studies B and C. The Study C figures were calculated by fitting a straight line to the capture times of each of the participants, and taking the mean of their times for capturing the first word (3.8 seconds) and for each additional word (0.91 seconds). Observed performance matched predicted performance closely throughout the critical range from 2 to 5 words, with less than 0.2 seconds error. The main error in the model appears to be an over-estimate of the time taken to capture additional words: our predicted time was 1 second, and our line fit gave 0.91 seconds. This difference is not significant given the small sample size, sample variance and measurement errors.

Table 2: Full results, including actual times using the auto-completion tool.

words	1	2	3	4	5	6	7	8	9	10
re-key, secs	2.0	4.0	6.0	8.0	10.0	12.0	14.0	16.0	18.0	20.0
prediction, secs	3.6	4.6	5.6	6.6	7.6	8.6	9.6	10.6	11.6	12.6
auto-comp, secs	3.8	4.7	5.6	6.5	7.5	8.4	9.3	10.2	11.1	12.0

5 DISCUSSION

The research described here has been interesting for us, both as a design exercise in its own right and as an indication of how critical parameters can potentially assist the design of a wide range of interactive systems. As a design exercise it generated some valuable outcomes:

- Our study led us to a critical parameter (short-segment text capture time) that we had not previously identified. Through further study this parameter could, we believe, be linked to overall writing speed.
- With re-keying times as a baseline, our designers were able to work with this critical parameter, and arrive at a satisfactorily innovative solution based on auto-completion.

- The parameter we identified proved to be manipulable, in the sense that performance predictions could be made that were accurate enough to guide the design team to an efficient solution.

The problem of designing a capture tool to support authors is far from solved. In particular, our solution assumes zero processing and recognition time, whereas in reality a time of 10 seconds or more is likely. We are confident this time can be reduced significantly. Now that its contribution to capture times (and possibly writing speeds) can be modelled, we hope that other researchers will tackle this problem as enthusiastically – and innovatively – as our team tackled the problem of rapid selection.

A question of wider significance is whether designing in terms of critical parameters, identified through studies of applications, can and should be practised more widely. There are two questions here: is it practical to design this way, and is it beneficial?

Use of critical parameters can become practical only if the parameters themselves can be identified. Our recent work confirms our earlier claim that critical parameters exist in many application domains [10]. We can now see that the derivation of predictive models is as important as the identification of parameters, and that the two combine to form a powerful basis for innovative design. However, our research also confirms that these parameters and models are time-consuming and often difficult to identify. This difficulty may persist; but if the parameters and models persist too then the effort will have been well spent.

Use of critical parameters and models will be beneficial if better interactive systems result. Our experience is that the results are indeed better: they provide better support to the user, and incorporate innovations that would probably not have emerged otherwise. There is a danger, of course, that parameters and models will be incorrectly identified, leading to performance 'improvements' and innovations that provide no help to the user. However, this can be ironed out by the iterative process of testing and validation. It seems likely that the use of critical parameters can lead to more real advances and useful innovations than are generated at present.

6 ACKNOWLEDGEMENTS

We wish to thank Ben Whitaker for his contributions to the design and construction of the auto-completion tool, and to the referees for their helpful comments on this paper.

7 REFERENCES

- [1] Alm N., Todman J., Elder L. and Newell A. F. (1993) "Computer Aided Conversation for Severely Impaired Non-Speaking People." *Proceedings of InterCHI '93 Human Factors in Computing Systems* (April 24-29, Amsterdam, Netherlands) ACM/SIGCHI, N.Y., pp. 236-241.
- [2] Byrne P. S. and Long B. E. L. (1984) *Doctors Talking to Patients*. Exeter: Royal College of General Practitioners.
- [3] 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.

- [4] Harper R. H. R. (1998) *Inside the IMF: An Ethnography of Documents, Technology and Organisational Action.* London: Academic Press.
- [5] Heath C. (1986) *Body Movement and Speech in Medical Interaction.* Cambridge: Cambridge Univ. Press.
- [6] Jaeger R. J. and Joel A. E., Jr. (1970) "TSPS No. 1: System Organization and Objectives." *Bell Sys. Tech. J.* Vol 49 (10), pp. 2417-2434.
- [7] John B. E. (1996) "TYPIST: A Theory of Performance in Skilled Typing." *Human Computer Interaction* Vol 11, pp 321-355.
- [8] 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.
- [9] Newman W. M. (1994) "A Preliminary Analysis of the Products of HCI Research, Based on Pro Forma Abstracts." *Proceedings of CHI '94 Human Factors in Computing Systems* (April 24-28, Boston, MA) ACM/SIGCHI, N.Y., pp. 278-284.
- [10] 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.
- [11] Newman W. M., Dance C. R., Taylor A. S., Taylor S. A., Taylor M. J. and Aldhous A. (1998) "CamWorks: A Video-based Tool for Efficient Capture from Paper Source Documents." *Proc. IEEE Multimedia Systems '99*, Vol 2, pp. 647-653.
- [12] Newman W. M. and Taylor A. S. (1999) "Critical Parameters in Interactive System Design: Pointers from a Literature-based Study of Representations of Performance." XRCE Technical Report EPC-1999-101, Xerox Research Centre Europe, 61 Regent Street, Cambridge CB2 1AB, UK.
- [13] Newman W. M. and Taylor A. S. (1999) "Towards a Methodology employing Critical Parameters to deliver Performance Improvements in Interactive Systems" *Human Computer Interaction: INTERACT '99*, pp. 605-612.
- [14] O'Hara K. and Sellen A. J. (1997) "A Comparison of Reading Paper and On-Line Documents." *Proceedings of CHI '97 Human Factors in Computing Systems* (March 22-27, Atlanta GA) ACM/SIGCHI, N.Y., pp. 335-342.
- [15] 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.
- [16] O'Hara K., Taylor A., Newman W. and Sellen A.J. (2000) "Understanding the Materiality of Writing while Reading from Multiple Sources." Submitted to *International Journal of Human-Computer Studies*.
- [17] Simon H. A. (1981) *The Sciences of the Artificial*, 2/e. Cambridge MA: MIT Press.
- [18] Vincenti W. G. (1990) *What Engineers Know and How They Know It: Analytical Studies from Aeronautical History*. Baltimore: Johns Hopkins Univ. Press.