

A Diary Study of Work-Related Reading: Design Implications for Digital Reading Devices

Annette Adler¹, Anuj Gujar¹, Beverly L. Harrison¹, Kenton O'Hara², Abigail Sellen²

¹Xerox Palo Alto Research Center
3333 Coyote Hill Rd.

Palo Alto, CA. 94304, USA

<adler, gujar, beverly>@parc.xerox.com

²Xerox Research Centre Europe

(formerly Rank Xerox Research Centre)

61 Regent St., Cambridge, CB2 1AB, UK

<ohara, sellen>@xrce.xerox.com

(Authors are listed alphabetically)

ABSTRACT

In this paper we describe a diary study of how people read in the course of their daily working lives. Fifteen people from a wide variety of professions were asked to log their daily document activity for a period of 5 consecutive working days. Using structured interviews, we analysed their reading activities in detail. We examine the range of reading activities that our subjects carried out, and then present findings relating to both common characteristics and variation across the sample. From these findings, we discuss some implications for the design of digital reading devices.

KEYWORDS

reading, writing, paper documents, digital document readers, design, diary study, field study, electronic books, virtual paper, user behaviour

INTRODUCTION

Technological advances in recent years have meant that digital display media are becoming more "paper-like". Wireless, lightweight digital displays are now more mobile and portable than ever. The quality and readability of the screens themselves are quickly approaching the brightness, resolution and contrast of plain paper. New techniques for stylus-based entry are making possible more paper-like interactions, such as allowing richer and more flexible mark-up and manipulation of digital documents.

Such advances have sparked a flurry of interest in developing digital devices for reading, as evidenced by a growing number of new projects using terms such as "portable document reader", "electronic book", "electronic encyclopaedia", "virtual paper", and "electronic paper". However, developers of these devices are faced with some difficult decisions. If they are to design devices for reading, what are the kinds of reading they should aim to support, who are the people that they should support, and how should they best support these reading activities? The fundamental problem is that the task of "reading" is far too general and ubiquitous: reading takes on a range of forms, is done for a variety of purposes, and is embedded within and related to many other document-based activities. These activities, of course,

will also vary depending on whether one is considering reading at work or at home, and will depend on the kind of work environment one considers.

The purpose of the study reported in this paper is to begin to answer these questions in a systematic, empirical way, within real work settings. While not discounting reading for leisure, this study focuses on work-related activities, whether they happen "at work" or in the home. We deliberately chose a sample of people who have very different types of jobs, across a range of different vertical markets. We did this because we were interested in exploring the range and diversity of work-related reading activities, as well as any commonalities that might nonetheless emerge despite this diversity. Our approach was to ask subjects to keep daily logs of their document activities, and then to unpack and expand these descriptions through structured interviews (a method successfully used previously, e.g., [2, 9]).

From this diary study comes a rich database of both descriptive materials as well as quantitative measures. In this paper, we present some selected findings addressing the following issues, which we feel have important implications for the design of digital devices for reading:

- A description and taxonomy of the range of reading activities and writing activities that occurred in the working lives of our sample.
- A discussion of common findings across our sample, most notably the predominance of reading in conjunction with writing, the predominance of paper-based document activity, and the degree to which reading activities occurred across multiple, independent reading surfaces.
- A discussion of the ways in which these people in different professions varied with respect to their reading activities.

Throughout these discussions we draw attention to what these findings mean in terms of design requirements for devices which aim to support on-line reading.

The Literature

To our knowledge there are no existing studies that attempt to characterise and quantify the range of reading activities that occur within real work settings, especially with an eye to the design of new technologies. There have been many attempts in the management science and business literature

Permission to make digital/hard copies of all or part of this material for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication and its date appear, and notice is given that copyright is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires specific permission and/or fee.

to analyse the various activities of office workers (see [4], for a review), but such analyses do not focus on reading activities, *per se*.

There is also a great deal of literature within psychology and human factors exploring the issue of how people read. Within psychology, the emphasis is mainly on modelling the reading process for specific kinds of reading tasks, and in strictly controlled laboratory situations. Having said that, there have been attempts to develop taxonomies of reading activity ([5, 6]). However, these are not derived from work settings and are more concerned with theoretical rather than practical implications of this categorisation. Within human factors, studies of reading do concern themselves with the practical application of reading, but much of the emphasis concerns on-line versus paper-based reading (e.g., see [1, 3]), again most frequently using short laboratory tasks testing psycho-perceptual parameters.

METHOD

Pilot Study

In order to develop a taxonomy of reading activities and to refine our methodology, we first ran a pilot study using 10 members of our respective laboratories, including ourselves, as "subjects". The initial taxonomy of reading activities was developed using a search of the literature [7]. Pilot subjects logged their document activities for 5 consecutive days (50 days worth of data). We then applied our taxonomy to each of these logs. All five co-authors discussed the data and the categories assigned, to produce consistency in the application of our taxonomy and some degree of inter-judge reliability. Based on these data, we refined the categorisation scheme, and developed a set of questions to ask our subjects in follow-up interviews. As a result of this pilot study, we also found it important to develop a taxonomy of writing activities in order to capture the range and extent to which writing occurred in conjunction with reading.

Choice of Subjects

This study was intended to be exploratory and, as such, we wished to gather broad and rich information about diverse reading activities that people carried out related to their work, across a wide variety of professions, representing a range of vertical markets.¹ In choosing subjects, we also took account of three factors for which we wanted a representative range:

- **Mobility:** We selected people whose jobs were predictably mobile across several locations (e.g., surgeon, anaesthetist, lawyer), unpredictable in terms of location (e.g., contract airline pilot, real estate agent), and generally centrally based in one location (e.g., accountancy assistant, warden).
- **Location:** We selected people from a variety of "activity locations" such as traditional offices, retail offices, home offices (i.e., self employed or working from home), and other work sites (e.g., cars, client sites, operating rooms).

¹ Clearly at some future point, deeper investigation of any one of these professions using many representative subjects would be of value. As it stands, with the possible exception of our cluster of medical professionals, this small and varied sample makes it impossible to generalise about any one professional group in particular. However, patterns across the whole sample are particularly suggestive given this diversity.

- **Collaboration:** We selected people who were different in the extent to which their work habits involved synchronous collaboration. Subjects worked as part of a tightly organised team (e.g., surgeon, organ transplant ward coordinator/nurse), loosely organised shifts (e.g., warden), or groups (e.g., lawyers). In other cases, subjects were primarily independent workers (e.g., optician, architect, real estate agent).

These criteria were partly derived from previous marketing analysis for portable document reading devices. They also represented issues the co-authors felt to be of critical importance to reading behaviour in this exploratory investigation. Based on these criteria, our final subject pool comprised: a contract airline pilot, a general surgeon, an organ transplant ward coordinator/nurse, two anaesthetists, a residential architect, a real estate agent, two lawyers (a director and an associate), a warden of a half-way house, an accountancy assistant, an optician, a marketing manager, an executive from a small start-up company, and a social worker.

Diary Method

Fifteen subjects were asked to log their daily document activities during the course of their working day, for five consecutive days. It was explained that by "document" we meant any sort of document, whether it be paper-based documents ranging from books to post-it notes, or electronic documents ranging from computer-based documents to the use of pagers or palmtops. We focused on documents that required the reading or writing of text. We asked subjects to log any document activity, not just reading activities. Further, we asked them to estimate the duration of each activity and to note this down on their log forms.

At the end of each day, an interviewer used a structured form in order to expand on the description of each activity and to note down more details such as what types of documents were used, whether the activity was collaborative, where the activity occurred, and what additional tools were used. Interviews typically took from 1/2 hour to 3 hours, depending upon the number of activities listed for the day.

By the end of the study we had 72 days worth of data across subjects. (The nurse worked only a 3-day shift, and the start-up executive 4 days that particular week.)

DATA ANALYSIS

The first step in the analysis was to apply our reading and writing taxonomies to the data. Classification of subjects' document activities into reading and writing categories relied, as much as possible, on subjects' own descriptions of their activities. However, in some cases, it was necessary to break down a single activity (as described by a subject) into multiple sub-tasks if the activity noted was discovered to involve multiple, distinct reading and writing components. As much as possible this was done within the context of the interviews and through discussion with the subjects.

Initial classification for both reading and writing activities was carried out by the interviewer of each subject. However, this was done with all experimenters/interviewers present so that ambiguous or difficult classifications could be discussed and agreed upon. Several sample logs were done jointly by all experimenters to ensure consistency in the

application of categories. Because this was an exploratory study with no specific *a priori* hypotheses in mind, we did not attempt to establish more formal inter-rater reliability using independent judges. We found that the reading and writing taxonomies could be unambiguously applied to most activities, though in some cases multiple categories needed to be applied to a single activity. Generally, the five experimenters felt that the categorisations were consistently applied, and any inconsistencies that did arise were "caught" and corrected in the course of the many subsequent analyses we carried out.

In order to obtain quantitative data on what proportion of their time subjects spent carrying out different reading or writing activities, the times associated with each activity involving reading and/or writing were analysed. Early in the course of the study we found it was difficult to obtain separate timing data for reading and writing within any given document-related activity. For example, a subject could tell us they spent 80 minutes reading and annotating a document, but not just how much of that time was spent reading as opposed to writing within that activity. Thus, we had little choice but to use the overall total duration of the document activity reported for each of the reading and writing sub-components. This means that the quantitative data we report will tend to overestimate the *actual* reading and writing times that occurred. However, the data reported do give an indication of the amount of time that people spent engaged in activities in which reading or writing activities either featured strongly, or acted in a supportive role.

RESULTS AND DISCUSSION

In general, if we assume a 40 hour work week for our subjects, we found that the document-related activity time they reported accounted for an average of approximately 82% of their working week, ranging from 23% (for the nurse) to 94% (the accountancy assistant). Such figures are comparable (though somewhat lower) to findings for knowledge workers in a "document-intensive" work setting [9]. Clearly, even in our diverse sample, document-related activities played an integral and central role in people's work, regardless of profession.

Reading in order to identify (ID) Glancing at a document only in order to identify which document it is.

Skimming (SK) Reading rapidly in order to establish a *rough idea of what is written*, and to decide whether anything is useful, or whether anything needs to be read in more detail later.

Reading own text to remind (REM) Reading *specifically* in order to remind oneself of what to do next, e.g., a To Do list, shopping list, Post-It note.

Reading to search/answer questions (SAQ) Reading to search for particular information: to answer a question, for reference, or to obtain information necessary to make a decision. *Goal-directed: the reader is sampling information in the text which satisfies the goal of the search.* Ranges from very simple goals to complex decision-making or problem-solving tasks.

Reading to self inform (SI) Reading for the purpose of furthering general knowledge *without any specific goal to which the information will be applied.*

Reading to learn (LE) Reading with the goal of being able to relate or apply information at a later date. This can refer to reading to review the basic concepts for discussion, or it can be much more reflective in nature.

Reading for cross-referencing (CR) Cross-referencing documents in order to integrate information. Includes reading from multiple sources or reading from one source. It can be for the purpose of writing, and may well include some editing activities.

Reading to edit or critically review text (ED) Reading in order to monitor what has been written in terms of content, style, grammar, syntax, and/or overall presentation. Includes editing one's own text, seeing how one's own text fits into a collaborative document, or the review of the text of others.

Reading to support listening (LI) Reading in order to support listening to someone else talk (e.g., following a presentation by looking at a series of slides).

Reading to support discussion (DI) Referring to text during a discussion in order to establish a mutual frame of reference and focus for discussion. Usually takes place in a synchronous collaborative situation (e.g., sitting around a table).

Figure 1. Categories of reading used in the analysis.

Reading Categorisation

Within subjects' document-related activities, reading occurred on average in 70% of them, pointing to the importance and also ubiquity of this kind of document activity. It is perhaps not surprising, then, that reading takes on many different forms within our data. The taxonomy of reading activities that we applied is described in Figure 1.²

The percent of time spent on document activities in which these reading activities occurred is shown in Figure 2, averaged across all 15 subjects. This includes activities where reading occurred by itself and where reading occurred in conjunction with writing.

As can be seen, the reading that occurred in our sample covers a broad spectrum of activity. For example, they range from "lightweight" sorts of reading activities such as quick glances to identify documents (ID), skimming (SK), and reading to remind (REM), to more intensive sorts of reading such as reading to learn (LE) or reading to edit text (ED). The reading activities reported also covered a range of different kinds of goals such as reading to search or answer questions (SAQ) or reading to self-inform (SI). They occurred in situations where people were working alone, as in reading to learn (LE), or in group situations, as in reading to supporting listening (LI) or discussion (DI).

Sketching out this range of reading activities has a general message for designers of digital reading devices which is that it is vital to understand precisely what is meant by supporting "reading". As varied as these categories are, the requirements for supporting one type of reading as opposed to another can be quite different. From Figure 2, we can see that for much of the work day, our subjects were involved in rapid and goal directed types of reading such as skimming (SK) and searching to answer questions (SAQ). This points to the need to consider digital reading devices that support fast and flexible search, manipulation and navigation. The predominance of cross-referencing activity points to the

² While our original taxonomy derived from the literature contained additional categories (e.g., reading in order to tell a story, reading for enjoyment, reading to follow instructions), the above list resulted from practical application to the work-related activities produced by our subjects. No doubt other categories might be necessary with a different sample of people, or in looking at reading for leisure, for example.

need for multi-document manipulation and viewing capabilities (CR). Finally, the high frequency of reading in support of discussion (DI) highlights the extent to which reading activity takes place in the context of other people using documents to support and provide a shared focus for that discussion. These kinds of reading can be contrasted with the simplified notion of a digital reading device as something which supports linear and continuous reading of single documents, such as books, by people who are alone. Certainly this does occur (as the categories SI and LE indicate), but the proportion of time spent in such activities is low by comparison.

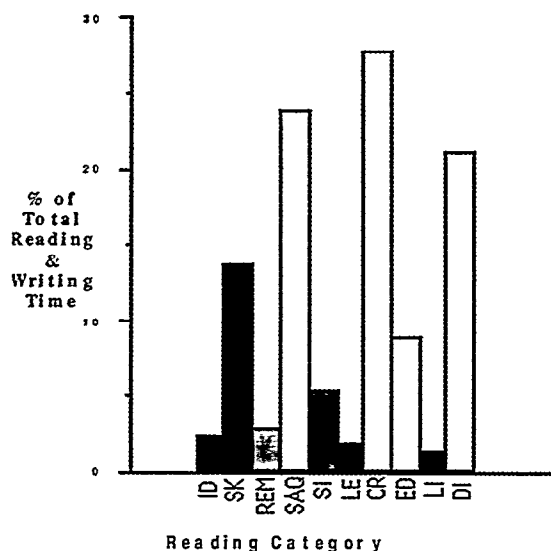


Figure 2. Reading Category Analysis

However, one may quite reasonably question the granularity of the above breakdown, especially if we are interested in specific design requirements for digital reading devices. It may be that this level of description is too broad for these purposes. To explore this issue further, we looked in more detail at the three categories of reading which occurred most frequently across our sample: the "skimming" (SK), "searching to answer questions" (SAQ), and the cross-referencing (CR) categories. Here, we will discuss only the first "sub-analysis" by way of illustration.

Analysis of Skimming (SK)

Within the skimming category (SK), we identified three sub-categories:

- (1) *Sorting* refers to skimming which typically occurred when subjects sorted through mail to determine what further action was needed, if any (e.g., reply, forward, throw away, file). This quick glancing to categorise documents also occurred when tidying up or filing away documents.
- (2) *Proofing* refers to a fast check of the contents of a document to ensure no items are missing or to quickly get the basic information out of a document. This more goal-directed kind of skimming includes examples such as glancing at medical records to familiarise oneself with content, reviewing patient histories, or looking over blue prints or flight plans, without actually searching for any particular item.

- (3) *Browsing* describes a sub-category of skimming which occurred mainly with newspapers, magazines, or journals where subjects said they did not read in detail but rather only to obtain a general sense of an article's content and extent. Typically this type of skimming appeared to have no immediate goal in mind to further apply what was being read (as do proofing or sorting).

Overall, no one kind of skimming dominated the data, with sorting occurring 38% of total time spent skimming, proofing occurring 31% of the time, browsing 28%, and "other" 3% of the time. What this analysis does indicate, though, is that categories within the reading taxonomy we have derived need to be examined more closely to understand the implications for design.

For example, in order for digital reading devices to support a sorting activity, one would require rapid and simple ways of moving through a "pile" of items or documents, such as showing the size of the electronic pile, with the topmost item always visible, and the possible destinations where items can be dropped into categorised piles. On the other hand, proofing typically occurs *within* a document. Devices supporting proofing require fast and simple page-by-page turning mechanisms. They should also provide a means of highlighting critical information *in situ* to preserve document context while giving the reader faster access to summary impressions of the document content. This is different from merely providing abstracts or summarised information separate from the context. Finally, the browsing sub-category occurs both within and across documents. Devices supporting this activity need to support both the rapid and simple page turning required within a document and a simple mechanism for switching documents. Browsing could take advantage of technology that highlights features of the articles, highlights critical segments, and/or provides brief summarisation. Functions similar to news clipping services might be used to provide overviews.

In summary, while the taxonomy that we applied to the data is useful in distinguishing amongst broadly different kinds of reading tasks, in order to use this categorisation scheme to help specify design requirements, some further sub-categorisation may well be necessary within each reading "type".

Commonalities across Subjects

Given the diversity of people, professions, and work settings represented in this study, and given the range of reading activities that emerged, it is instructive to ask whether these reading activities exhibited any common characteristics, despite the underlying variation, from which we might draw some general design implications. Here we highlight three issues in particular: the extent to which reading (and writing) was paper-based, the degree to which reading took place in conjunction with writing, and the use of multiple "display" surfaces in the course of reading and writing.

Paper Versus On-Line Reading and Writing

In order to examine the extent to which reading and writing activities were paper-based as opposed to being carried out using on-line tools, we classified each reading or writing activity as falling into one of three categories: paper-based, on-line, or presentation-based. *Paper-based* reading and writing included such items as drawings, blue prints, flight plans. *On-line* reading and writing tended to be computer-

based but included specialised equipment with screens such as medical monitors and weather machines for flight reports. In this category we also included situations where subjects moved between on-line and paper media. Third, we separated out *presentation* type media which included videos, 35mm slides, and overheads.

Analysis across all subjects showed that paper-based reading and writing accounted for 85% of people's total activity time. On-line reading and writing make up about 13% of the total time, and presentations only 2%.

These figures are especially interesting in light of the fact that all of our subjects had computers with essential data which required them to use the computers for at least some portion of their work. One can surmise that at least in some situations, the choice of paper over an available on-line alternative is because paper supports their particular reading and writing tasks better, and that the on-line alternatives simply fail to provide the same critical affordances as paper [8], [9]. The next two findings, the degree to which people write as they read, and the use of multiple documents during reading, provide some insight into this issue.

Reading in Conjunction with Writing

Early on in our analysis, it soon became clear that, within a given document activity, reading frequently co-occurred with some kind of writing. To assess the extent to which this occurred, we tabulated the total time for activities in which reading alone occurred (without writing), as against the total time in which reading occurred with writing. We found that, while there is significant variation, reading was accompanied by writing most of the time. In fact, for 8 of the 15 subjects, reading occurred with writing over 75% of the time and up to 91% of the time. For all but one subject, reading was more often accompanied by writing than not. In order to better understand what kinds of writing were occurring in our sample, we applied the taxonomy shown in Figure 3.

Creation	Writing to create a new document or writing to modify an existing document.
Note-taking	Writing in an abbreviated and unstructured way, usually in order to serve an intermediate or temporary purpose rather than to create a final document.
Annotation	Marking on top of an existing document to relate the marks to their surrounding context.
Form-filling	Filling in structured forms or writing in a prescribed, structured way, such as filling in a register.
Updating	Updating calendars or schedules.

Figure 3. Categories of writing used in the analysis.

One design implication which comes from this analysis is that if one is to design a digital reading device, it makes little sense to think of reading in isolation from writing.

The breakdown of writing activities occurring in the context of reading, averaged over all subjects, is shown in Table 1.

Writing Category	% Total Reading & Writing Time
Creation	16.4%
Note-taking	21.7%
Annotation	26.4%
Form-filling	33.9%
Updating	1.6%

Table 1. Percentage of different writing activities (done in conjunction with reading), averaged over all subjects.

As shown in Table 1, writing includes far more than the creation or modification of stand-alone documents. It appears that a significant amount of time, at least in our sample, was spent form-filling while reading, or annotating documents while reading them, or taking notes in the course of reading. Thus it appears that a digital reading device which does not support what is mostly pen-based marking either in conjunction with, or on top of documents, will have limited value if one is seeking to support a wide range of work-related reading activities.

Cross-Document Use

The final common finding that emerged that we will discuss is the use of multiple documents during reading. As is indicated by the dominance of the reading category "cross-referencing", it appeared that subjects were carrying out a great deal of their reading and writing activities using multiple documents in parallel with each other.

In order to examine this issue more closely, all of the subjects' reading and/or writing activities were categorised as using a single surface for reading and/or writing or whether multiple display surfaces were in use concurrently for reading and/or writing activities. In this analysis, a display surface was defined as an independent, tangible surface such as a computer screen or piece of paper. In order for them to be classified as used concurrently, it had to be clear that the surfaces were being used side-by-side for the same activity. Here, flipping back and forth within a document was not counted as a case of parallel, concurrent usage. Multiple windows on a computer were considered a special case of page flipping and not concurrent usage. (Windows usually needed to be switched around to bring any one into the foreground to be read, much like page flipping.)

When analysed on the basis of total document-related activity time, we found that 52% of the time subjects used a single surface for reading, writing, or a combination of reading and writing, such as annotating a document while reading it. This means that 48% of the time they used *at least two display surfaces* concurrently. To explore this cross-document activity in more detail, we further categorised these activities as follows:

- *Independent Reading and Writing Displays*: Activities in which one display was used primarily for writing while cross-referring to or reading from another independent display.
- *Multiple Reading Displays*: Activities in which no writing occurred but in which two or more displays were read with reference to each other in parallel in the same activity.
- *Multiple Reading Displays plus Writing*: Activities in which two or more displays were used for reading with

reference to each other, as described above, but in which a writing activity also took place, either on a display also used for reading or on a separate display.

The breakdown of this cross-document use for all 15 subjects is shown in Table 2. As can be seen, the majority of cross-document activity involved reading in combination with writing, and the biggest portion of that involved the reading of more than one document. In these cases, writing either occurred on a separate document or occurred on one of the documents also being read.

Multiple Document Category	% Total Reading & Writing Time
Independent reading & writing surfaces	19.1%
Multiple surfaces for reading only	3.9%
Multiple surfaces for reading plus writing	25.1%

Table 2. Breakdown of use of multiple display surfaces as a percentage of total activity time involving reading or writing.

Just why these sorts of activities were so frequent is best explained by looking in more detail at what subjects were doing when reading and writing using multiple documents. This can be summarised as follows:

- *Extracting information.* One reason separate notes or form-filling occurred while cross-referring was to extract information from the document being read - as a list of reminders for themselves, to draw attention to some subset of the information in the document being read, or to instruct others to do things. Sometimes information was extracted in order to keep a separate record of it. Extracting information could also involve some degree of transformation of that information, such as using a separate piece of paper as a computational workspace, or for writing summaries, or plans of action.
- *Integrating information.* Another major reason for cross document activity was to integrate information from other sources. This could be in the service of constructing another document. Subjects referred across documents both for specific facts and figures, and also for general ideas and themes. Cross referencing was also done when updating one document using another, such as updating figures in a database, incorporating paper-based changes into an electronic version, or using old "To Do" lists to create new ones. Finally, information integration across documents was done to make decisions, by considering information from multiple different sources.
- *Consistency checking.* A third reason for cross document activity was to ensure that one document was consistent with another. This could involve checking for consistency amongst facts and figures, or could involve checking for consistency in language or format.
- *Critiquing or making comment.* Finally, cross document activity also took place when subjects had to critique, comment or refer to a document in order to write another.

One clear implication from this is that the sorts of activities described above somehow need to be supported in digital reading devices if they are to satisfactorily replace current practice. Either this needs to be accomplished by providing at least two readable/markable display surfaces, or the ability to carry out these various activities has to be supported in other ways. For example, one could imagine, and indeed

there already exist, on-line facilities for ensuring consistency (i.e., templates ensure format consistency). Interfaces are being developed which accomplish some of the kinds of information extraction tasks subjects reported doing with multiple documents. For example, the Dynamite notebook [10, 11, 12] allows certain kinds of information such as action items and reminders to be indexed and automatically compiled separately from the document in which they are embedded. However, it may be that the multiple display, "paper-like" approach is, in the final analysis, a much more parsimonious approach to supporting the whole range of tasks that subjects accomplish by using more than one document.

Variation Across Subjects

Up to this point we have been focusing on commonalities that emerged across our subject sample, and some general design implication. Of course in such a varied collection of people, across such a wide range of professions, one would expect interesting differences in the ways in which they read, in the ways in which they combine reading with writing, and in other aspects of where and how reading occurs. These are especially important to understand if, as a designer, one wishes to target a specific set of "users" in a particular vertical market. Of course, one cannot make such generalisations about the needs of an entire profession or vertical market on the basis of our small sample, but it is nonetheless interesting to look more closely at the inter-subject variation that did occur to get some sense of how different people's needs might be. While we are forced to be brief in our discussion of these issues, we will concentrate on four aspects of that variation: the "top" reading and writing categories for each subject, degree of sharing, and mobility.

Table 3 shows, for each subject, which reading and writing categories featured most strongly. Specifically, it considers which of the reading and writing categories constituted the highest percentages of total reading and writing activity time for each subject.

Table 3 also shows the extent to which each subject tended to engage in shared document viewing. We considered the shared viewing of a document either to be a single document which two or more people gather around to see simultaneously or multiple copies of a single document which people read simultaneously as a reference. Here, sharing involved either the joint reading of a document (such as when people are reviewing case histories or medical charts together), the joint writing of a document (such as where several people may be viewing a document and dictating modifications or additions to one person), or combined reading and writing such as is typical of joint authorship situations where a document is reviewed and annotated in a collaborative setting.

From this analysis we found that subjects fell into 3 different clusters: "high" sharers were those where shared viewing occurred for over 50% of their activities; "medium" cases were those people in which shared viewing occurred for about 50% of their activities; and "low" sharers were those for whom shared viewing occurred for less than 50% (and in this case less than 30%) of their activities.

The final dimension represented in Table 3 is that of "mobility". Here, all reading activities were classified according to where they took place: in the subject's office (including offices located in subjects' homes), within the site where the subject's office was located, or outside of the building or site. We found that subjects could then be grouped on the basis of the amount of time spent on reading activities falling into these three categories. Nine subjects fell into the "non-mobile" group, meaning that the largest percentage of time spent in reading activities was done at the desk in the office; two subjects were primarily "locally mobile" in that the largest percentage of their activities were within the office site but outside of their office; and two subjects were mostly "remotely mobile" in carrying out most of their activities involving reading outside of the site containing their office. Two further subjects we classified as "combination" cases, where approximately half of their time was spent at their desk and half was spent being remotely mobile.

Subject	Top Reading Category	Top Writing Category	Shared Viewing Category	Mobility Category
Architect	Edit/Review	Creation	low	combinat'n
Pilot	Cross-Rcf	Note-take	medium	remote
Warden	Cross-Ref	Form-Fill	medium	non-mobile
Account. Ass't	Cross-Ref	Form-Fill	low	non-mobile
Lawyer (Assoc)	Cross-Ref	Creation	low	non-mobile
Lawyer (Dir.)	Discussion	Annotate	medium	non-mobile
Optician	Discussion	Form-Fill	high	non-mobile
Social Worker	Discussion	Note-take	high	non-mobile
Market'g Mgr	Discussion	Annotate	high	combinat'n
Start-Up Exec.	Discussion	Creation	medium	remote
Real Est. Agent	Search/Ans	Note-take	low	non-mobile
Nurse	Search/Ans	Form-Fill	low	non-mobile
Doctor	Search/Ans	Form-Fill	low	non-mobile
Anaesthetist 1	Search/Ans	Form-Fill	low	local
Anaesthetist 2	Search/Ans	Form-Fill	low	local

Table 3. Four dimensions of inter-subject variability.

The first interesting feature of this analysis is to note the similarities amongst the four medical personnel in our sample (doctor, 2 anaesthetists, and nurse). All can be characterised as form-fillers, whose main reading task is one of searching through documents for facts, results, instructions, and other specific kinds of information, often in order to complete those forms. This kind of document activity is something they do on their own (although the documents they interact with are generally jointly authored in themselves). Note, however, that these medical professionals differ in the degree to which they are based in an office, with the anaesthetists generally using public or shared spaces in which to carry out their document work.

A second interesting group of subjects is those who spend a great deal of time reading in support of discussion (lawyer-director, optician, social worker, marketing manager and start-up executive). As one might well expect, these people are also in the "high" or "medium" shared viewing cate-

ries. Having said that, they vary in their mobility, and engage in very different kinds of writing activities. For example, the optician is mainly form-filling as a result of discussion with clients, while the social worker mainly took free-hand notes. The marketing manager relied very heavily on the annotation of the documents around which the discussions were centred in order to record pertinent facts, or to write reminders for herself.

The third "group" which emerges is what one might call the "cross-referencers". However, on closer inspection of their data, the nature of the cross-referencing these subjects carried out was actually quite different from subject to subject. For example, the accountancy assistant spent most of her time cross-referencing to check for consistency amongst figures, the pilot was integrating and extracting from multiple documents to plan routes and check flight information, the warden was mainly extracting information for form-filling, and the lawyer-associate was mostly cross-referencing to check language and format of documents in the creation of new ones. While one might surmise that such heavy use of multiple documents is best suited to an office environment where papers can be spread out in space, obviously remotely mobile workers engage in this sort of activity too, as the case of the pilot demonstrates.

The first implication of this overview of inter-subject differences is that one might conceptualise a range of different kinds of portable, digital reading devices. For example, some might be specifically geared toward providing users with an armoury of different kinds of flexible search and navigation features for people whose main reading task is that of quickly browsing through information, or finding and getting access to specific kinds of information. Digital devices obviously have the potential to confer many benefits here, especially in terms of increased access to more information, and to more up-to-date information. At the risk of over-generalising, it appears that medical practitioners are one group of users who might benefit from this sort of device. For example, in medical settings, paper documents are often problematic in that patient's notes are often out of date, information needs to come from many different departments, and forms get lost. Digital reading devices specifically designed to allow quick, flexible access to information could have an important impact here, especially if careful consideration is given to the ways in which medical practitioners currently search through documents.

Another "kind" of digital reading device is one that is designed with collaborative use in mind. This will affect both hardware and software requirements. For example, the viewing angle of the screen will be an issue here. One can imagine exploring a range of possibilities for allowing multiple pen input, and for providing software which supports joint marking and editing. Our data suggest that such a device might be used across a range of different job types.

The group of cross-referencers suggests that some people may be more heavily in need of multiple display surfaces than others in their work. Thus a "two-screen" version of a device might well find its niche naturally within some professions. However, judging by the degree to which multiple displays were used across our whole sample, we can surmise that at least two functionally interlinked screens might

be more generally useful, for the whole range of reasons we have outlined.

As a final comment on inter-subject variation, the data on the amount of reading carried out away from the desk strengthens the case for the need for truly mobile reading devices (see also [3] for empirical results confirming the dimensions of this). Even for the 9 "non-mobile" subjects, 6 of these people spent from 13-40% of their time doing reading activities away from their home office. (The optician, warden, and social worker were entirely desk-based.) Obviously there are a range of reasons why people are office-based, and we can only guess at the extent to which these figures would change if they had technologies available which enabled them remote access to the document tools and resources currently only found in their offices.

CONCLUSION

We have shown how a diary study can allow the collection of both descriptive and quantitative data on reading activities within the context of people's working lives. Further, we have used our analysis to draw out several different design implications for the design of new technologies in support of reading. These can be summarised as follows:

- The linear, continuous reading of single documents by people on their own is an unrealistic characterisation of how people read in the course of their daily work. Reading takes on a variety of different forms, serves a range of different goals, and is carried out in many different ways, as the 10 reading categories we have identified illustrate.
- Looking in detail at the way reading is carried out within any one reading category can offer guidance in specifying a range of design requirements.
- Reading occurs more frequently in conjunction with writing than it does in isolation. Thus, it appears that writing (in a variety of forms) is an integral part of work-related reading. Designers need to seriously consider the importance of developing reading devices which also support the marking or writing of documents during the reading process.
- Almost half of all the document activities which involved reading involved the concurrent use of multiple "display surfaces" for reading, or reading in conjunction with writing. This points to the need to consider how single display devices can support the range of cross-document activities people carry out. It also emphasises the importance of considering the benefits of designing devices which are based around the use of multiple, digital displays.
- Inter-subject variation suggests the potential for a range of kinds of digital reading devices. Examples include a reading device specifically developed to support goal directed searching and browsing (such as was characteristic of our medical personnel), and a reading device designed for collaborative use.

In drawing attention to these issues, we hope that we have illustrated the value of studying reading behaviour in detail, and in the context of day to day working life. While guesswork and intuition can result in very successful design solutions, we suggest that grounding design in empirical data

can be a productive, more systematic way forward. At the very least, it can help us to clarify our assumptions, and ask important questions about who we are designing for, and what is the nature of the current working practices we are aiming to augment or change for the better.

ACKNOWLEDGEMENTS

We are grateful to all of our subjects for taking part in this study and allowing the intrusion into their daily working lives. We also wish to thank Marge Eldridge for comments on this paper and assistance in setting up the study, Joe Sullivan and his group at FX-PAL for providing additional financial and moral support, and Tom Moran for comments on the paper, input on the study design, advice and support.

REFERENCES

1. Dillon, A. (1992). Reading from paper versus screens: A critical review of the empirical literature. *Ergonomics*, Vol. 35(10), 1297-1326.
2. Eldridge, M. & Newman, W. (1996). Agenda benders: Modelling the disruptions caused by technology failures in the workplace. *Proceedings of CHI '96, Companion Vol.* New York: ACM Press, 219-220.
3. Gujar, A., Harrison, B. L., and Fishkin. (1997). A Comparative Empirical Evaluation of Display Technologies for Reading. Submitted to CHI'98 papers.
4. Hirschheim, R.A. (1985). *Office Automation: A Social and Organisational Perspective*. John Wiley and Sons, Chichester, New York.
5. Kintsch, W. and van Dijk, T.A. (1978). Toward a model of text comprehension and production. *Psychological Review*, 85, 363-394.
6. Lorch Jr., R.F., Lorch, E.P. & Klusewitz, M.A. (1993). College students' conditional knowledge about reading. *Journal of Educational Psychology*, 85, 239-252.
7. O'Hara (1996). *Toward a typology of reading goals*. XRCE Technical Report No. EPC-1996-107. Xerox Research Centre Europe, Cambridge, UK.
8. O'Hara, K., & Sellen, A.J. (1997). A comparison of reading on-line and paper documents. *Proceedings of CHI '97*. New York: ACM Press.
9. Sellen, A.J., & Harper, R.H.R. (1997). Paper as an analytic resource for the design of new technologies. *Proceedings of CHI '97*. New York: ACM Press.
10. Shilit, B., Golovinchinsky, G. and Price, M. (1997). Beyond Paper: Supporting Active Reading with Free-form Digital Ink Annotations. Submitted to CHI'98 papers.
11. Wilcox, L., Shilit, B., and Sawhney, N. (1997). Dynamite: A dynamically organised ink and audio notebook. *Proceedings of CHI '97*. New York: ACM Press, 186-193.
12. Bly, S., Wilcox, L., Chui, P., and Gwizdka, J. (1997). Finding Information in Ink and Audio Notes: A Study of Indexing. Submitted to CHI'98 papers.