City Research Online

City, University of London Institutional Repository

The Digital Reading Desk: A lightweight approach to digital note-taking ☆

Jennifer Pearson[a], George Buchanan[b,*], Harold Thimbleby[a], Matt Jones[a]

[a] Future Interaction Technology Laboratory, Computer Science Department, Swansea University, Swansea SA2 7BP, UK

[b] Centre for HCI Design, 2nd Floor College Building, City University, Northampton Square, London, EC1V 0HB, UK

ARTICLE INFO

ABSTRACT

Attentive reading is a complex and cognitively demanding task that uses note-taking and annotation to support the reader's interpretation of the document. When reading on paper, extensive use of highlighting and other activities are conducted to support attentive reading, but this rich behaviour is not used with digital documents. Many users therefore print digital documents and then interact with them in physical form.

This paper presents the "Digital Reading Desk," an enhanced digital reading environment that provides support for attentive reading, providing a large working space for notes that mimics the use of desk space in a conventional physical environment. The Reading Desk uses a single tool to support both annotation and bookmarking, simplifying both the user's learning of the system and their use of tools. Evaluation of the Reading Desk indicates preference for it, as well as considerable behavioural differences between our approach and traditional digital reading tools.

© 2012 British Informatics Society Limited. Published by Elsevier B.V. All rights reserved.

1. Introduction

Reading is an ubiquitous human activity that is increasingly being performed interactively on-screen, whether using traditional PCs or mobile reader devices such as the Amazon Kindle or Apple iPad. Studies of the use of digital text have repeatedly demonstrated that poor human–computer interaction in digital reading inhibits user performance (Marshall and Bly, 2005; O'Hara and Sellen, 1997). Whilst research to improve the performance of these reading tools has been on-going for some time, progress has been slow. The increasing ubiquity of digital reading appliances—such as tablets and eInk devices, which mostly replicate the interaction design of current desktop reader software—means that it is timely to scrutinise the design and usability of electronic books and articles.

Attentive or "active" reading (Adler and Doren, 1972) is a complex and multi-faceted activity that takes years to master. Previous research has indicated that printed text on paper is a more effective medium for this work than digital documents (O'Hara et al., 1998). However, the underlying causes of this difference are not perfectly understood, and both the causes of digital inferiority and the design of improved digital reading tools are on-going areas of research (Marshall and Bly, 2005).

One concept that provides a theoretical underpinning of effective active reading tools is what Marshall et al. describe as *lightweight navigation*: that is, "navigation that occurs either when people reach a particular page or when they move within an article in a way that is so unselfconscious that they are not apt to remember it later" (Marshall and Bly, 2005). This suggests that navigation, and related tool use, is or should be a secondary task for the user and receives little conscious attention. This idea of "lightweight" interaction is supported by Csikszentmihalyi's theory of flow (Csikszentmihihalyi, 1990), which describes a state of mind where people are so involved in an activity that, effectively, nothing else matters. Brought together, these two concepts suggest that the tools used to aid active reading should require minimal attention to maximise the cognitive resources available for the main reading task. We can thus conclude, that in order to support close attentive reading, a digital reading environment should provide tools that place low demands on the user's cognition.

On paper, this ideal is often realised: annotation can be so subconscious, that the user is unlikely to remember doing it. The physical properties of paper afford many actions (e.g., folding and flicking) that all contribute to the undemanding manipulation of the printed document. In contrast, the analogous digital tools are seldom used during reading (Sellen and Harper, 2003): while acquiring reading skill requires considerable time, computer-based tools are usually either quickly adopted or discarded. The fact that basic tools of digital reader software are not used, demonstrates that there is clear scope for improvement.

We present a novel digital reading system that aims to support the needs of users engaged in attentive reading. Our design intends

to reduce the cognitive attention that supporting tools require from the reader. It uses a single form of note that can be used for both bookmarking and annotation, and we extend the workspace of the user beyond the logical pages of the document: in both cases replicating lightweight properties from paper. The system we report here is designed for a desktop PC environment, which remains the primary digital hardware for active reading on longer texts (Czerwinski et al., 2003).

This article extends work reported in our previous short paper (Pearson et al., 2011) and reports a fuller range of results and findings in greater detail. We start with a review of current research on place-holding and annotation techniques. After motivating the project and summarising previous research, we report the design and implementation of our Digital Reading Desk system, which is then followed by the description of a comparative user study that assesses the system, and we conclude with a short discussion of ideas and possible areas of future study.

1.1. Motivation

Research has consistently demonstrated that current digital note-taking tools suffer from poor interaction design (O'Hara and Sellen, 1997; Schilit et al., 1998a; Marshall, 1997). Investigating how users interact with physical documents provides a useful foundation for understanding the shortcomings of digital document interaction. Adler et al. (1998), focus on how users undertake work-related reading, and they identify the types of reading activities, purposes and readers that software should support. They found that reading more often occurs in conjunction with writing than in isolation. The authors therefore suggest that any digital reading device which does not support marking will have limited value in work-related reading activities. The paper reports that users construct a contextual understanding of the text as they read, and that writing and drawing (e.g., underlining) supports their primary reading activity.

O'Hara and Sellen's study of electronic document use (O'Hara and Sellen, 1997) demonstrates that although digital documents offer some advantages over paper (e.g., search), in the case of reading *for the purpose of* (later) writing, the benefits of paper far outweigh those of online documents. They also establish that the causes of this are not limited to, and indeed are minimally influenced by, physical display properties such as screen resolution or contrast. Rather, they argue that that the differences in paper and online documents are more to do with the advantages that paper offers in annotation, fast navigation and spacial layout. The authors conclude that understanding the use of paper can help improve the design of digital reading technologies.

1.2. Problems

Current digital note-taking tools lack the usability of their physical paper counterparts, which results in low rates of use and many users opting to print documents to read them (Sellen and Harper, 2003). As we move closer to the ideal of a paperless office (Sellen and Harper, 2003), it becomes more important that document mark-up tools are as easy to use as possible. Our goal is to further understand how to provide truly lightweight interaction by building on evidence from earlier research (Buchanan and Pearson, 2008; Pearson et al., 2009; Schilit et al., 1998b). Before we continue to the design of our implementation, we briefly describe three problem areas we believe contribute to the unintuitive nature of digital readers.

1.2.1. Space

The result of recent studies on paper and digital annotation (Pearson et al., 2009) brought us to the conclusion that the margins, and the space surrounding the document (e.g., the desk), perform an integral role in the physical mark-up process. Fig. 1 shows an example of note-taking on paper: here, the desk area surrounding the book is being used to keep notes that remain in place regardless of the currently open page. Unfortunately, there is no comparable workspace in most digital interfaces, a factor we believe contributes to the poor usability of digital note-taking tools. See Aim A1, in Section 3.1.

1.2.2. Multiple complex tools

On paper, it is common for one tool to have multiple functions. For example, a ruler can be used to measure as well as to draw straight lines (Dix, 2007). Similarly, physical pages can be marked easily by simply slipping in a piece of scrap paper or sticking a Post-it. Such placeholders can also serve as a note simply by being visible, and of course by any writing on it (Fig. 1). Digitally however, these functions are typically separated into two distinct tools: such as the tree list bookmark structure of Adobe Acrobat (Fig. 2) and its separate note-taking facility. Previous work (Buchanan and Pearson, 2008) on place-holding has used a visual interface with coloured "tabs" for bookmarks (Fig. 3), giving more visibility than existing interaction designs but still does not allow multiple uses. See Aim A2, in Section 3.1, for details on how we plan to overcome the problem of tool overload within digital documents.

1.2.3. Menu navigation

When a user is engaged in active reading, it is vital that they devote as much time as possible to their main task, and not to the control of supporting tools. Csikszentmihalyi's theory of "flow" (Csikszentmihihalyi, 1990) describes a state of complete concentration upon a task, to the point at which a user is so absorbed that all other considerations are forgotten. To experience flow, the user must be highly skilled and engaged in a challenging task, without distractions that divert the focus of attention. We hypothesise that annotation and bookmarking tools are secondary to the main task of "reading." Indeed, researchers such as Adler et al. (1998) and Marshall and Bly (2005) have suggested that the conscious effort demanded by digital note-taking tools impedes the main active reading task: a view corroborated by the theory of flow. Therefore, our intention is to arrive at an interaction where the tools become transparent to the central active reading task, where flow should be experienced.

Turning the user's attention away from the main text to control features such as menus will, likely, displace their focus from the text. We exploit direct manipulation (Shneiderman and Plaisant,
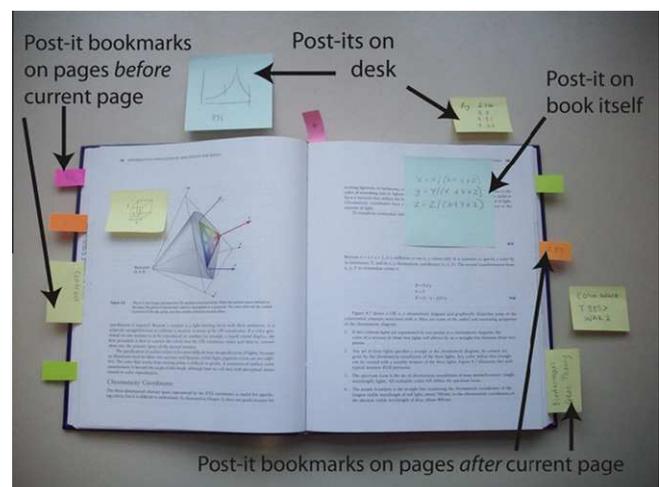

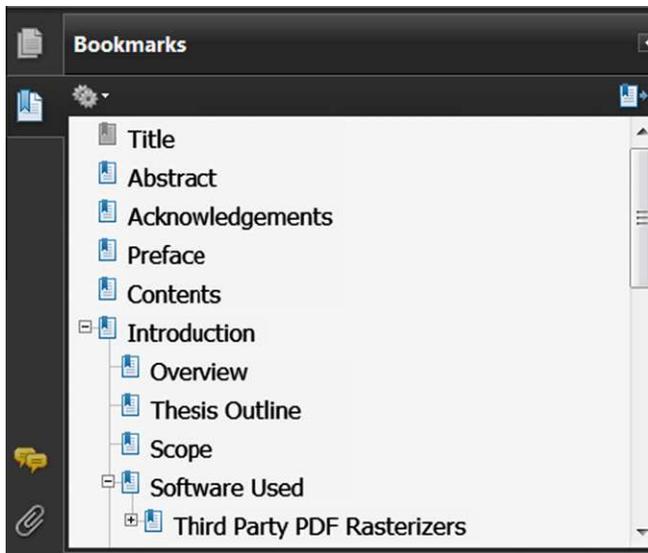**Fig. 1.** An example of Post-its being used in a physical book.

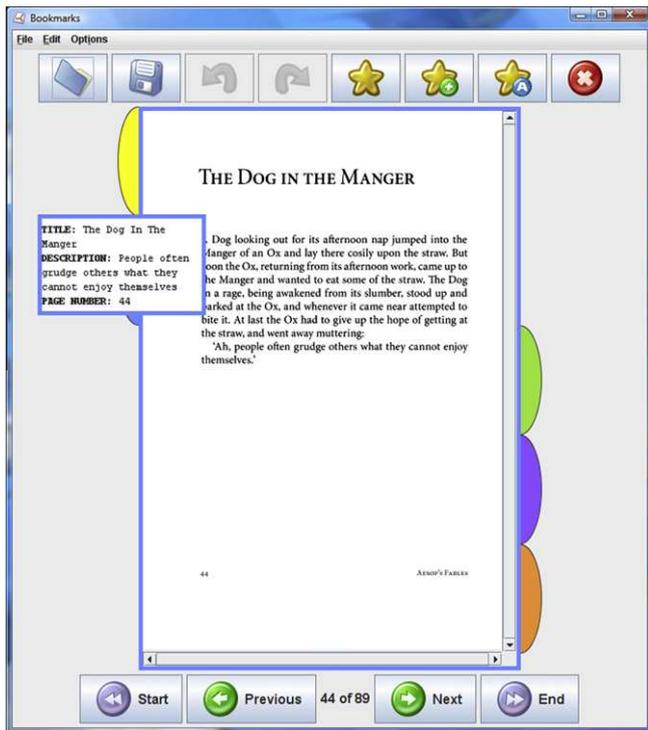**Fig. 2.** Example of the tree bookmark structure used in Adobe Acrobat.

Fig. 3. The Visual Bookmarking system. The coloured 'tabs' represent bookmarked pages: those that appear on the top left are on pages that are chronologically *before* the current page and those that are on the bottom right are chronologically *after* the current page. Changing the current page will cause the bookmarks to "flip" accordingly.

2004) to minimise the number of menu controls, and consequently maintain the user's visual focus close to the user's point of attention in text. See Aim A3, in Section 3.1.

2. Related work

This paper focuses upon user interaction issues in the creation and editing of digital notes. This general area and topic has been studied by researchers from a number of different disciplines over recent years. Digital library researchers have investigated both technical issues for storing and retrieving annotations (e.g., Frommholz and Fuhr, 2006), and the human–computer interaction issues surrounding note-taking work (Marshall, 1997). In this section we investigate the existing research on the HCI of reading and note-taking, beginning with broader socio-technical issues, before turning to concerns of interaction design.

2.1. Reading in digital and print media

The comparison of physical and digital reading has been an ongoing area of research for over thirty years. Early research investigated fundamentals such as comprehension, speed and efficiency (O'Hara and Sellen, 1997; Dillon, 1992). Later research embarked on a wider range of issues, and established reading as a complex activity that required significant support for the user's cognition. The majority of this research has rated the performance of digital documents unfavourably compared to paper (Sellen and Harper, 2003).

Annotation has been compared between digital and print media (Sellen and Harper, 2003; Marshall, 1997). The outcome of these different studies has been a growing body of evidence that there is a lower rates of annotation on conventional digital documents than on paper. This mirrors the differences observed in studies of reading in general. There is general pattern of behaviour where users prefer to print and annotate or mark-up the paper copy of an electronic document (Sellen and Harper, 2003; Marshall, 1997). Investigations of reading software design drew similar conclusions, and reinforced the building evidence that annotation is a critical support to a user's attentive reading (Schilit et al., 1998a).

In addition, research that studied reading within the context of its physical environment demonstrated that people use this physical space to organise their work and place meaning on it (Adler et al., 1998), another pattern confirmed when reading software was evaluated (Marshall et al., 1999). This result is consistent with the findings of key literature that shaped the field of spatial hypertext (e.g., Malone, 1983).

2.2. The book metaphor

Given the apparent advantage of print, some researchers have sought to reproduce the experience of physical books as literally as possible when using digital texts. This allows the direct transfer of the behaviours from the physical to digital domain. Different approaches have been taken within this general paradigm. We first examine research into "realistic books" that present a very literal, visualisation of a text as if it were a "real book." We subsequently study the ways in which printed books have been supplemented with digital technologies to enhance the user's experience of reading.

2.2.1. Realistic books

Liesaputra and Ian (2008) produce a visual reproduction of a printed books on the computer display. The "Realistic Book" system closely imitates paper, including a double-page spread, pages curling as the user navigates between them, and applying ageing processes to heavily read pages. The paper-like behaviours are primarily visual, rather than interactive (e.g., the appearance of pages being turned). The closest similarity to our work appears where bookmarks project beyond the physical paper. However, the position and size of these are standardised and computer-controlled, as opposed to our freeform, user-controlled approach.

Later work expanded on this concept, with modifications to the rendering of the "book" and the addition of new features such as search, and enhanced facilities, including bookmarks. These often combine; for example, if a search is done across a book, marks are displayed that indicate the each page which contains a hit.

These tags appear in much the way that Post-it tags would if they projected out from the page (Liesaputra et al., 2009).

However, there have been limitations to these literal approaches. The turning of the page, whilst aesthetically pleasing, is not necessarily a positive contribution to usability. The response speed of the program is reduced, due to the computational cost of rendering page turning in detail. Besides such detail, the behaviour of the programs is much akin to previous document reader software. The use of tabbed markers to indicate search hits was not closely examined by experiment, and thus the potential of that approach is not scientifically established. While the overall performance of the realistic book—compared to a PDF—was established in terms of time-efficiency, which elements of the design were critical to achieving this was not made clear. It is also not clear how other uses of the same rendering technique would perform, in terms of usability.

2.2.2. Augmenting real books

Wu et al. (2008) supplement physical books with digital, interactive services. This method provide digital enhancements, where there is perceived to be an electronic advantage—e.g., in providing hyperlinked material, projected around a physical book. This strand of research frequently exploits the perceived value of the working space that contextualises physical items, and places the individual text within a richer reading environment that expands beyond the boundaries of the page. Whilst we do not follow their interest in physical books as the reading medium, we exploit the "space" around a document to help enhance the user's active reading. The key difference is that our interest is in digital documents.

2.2.3. Reading appliances

A final way in which the metaphor of the book has been explored is the use of so-called "reading appliances." The foremost project in this area is XLibris, which created a tablet-like device that supported reading different texts and provided a variety of annotation tools. The emphasis of the project was on studying reading outside of the desktop PC, and the impact of both physical form and interaction on the overall reading experience.

Our own research endeavours to be device independent; therefore some of our concerns diverge from the appliance focus of XLibris. Nonetheless, the XLibris project did investigate user annotation practices, including the use of free-form ink annotation (Schilit et al., 1998a). Our own focus here complements the XLibris work: our interest in space is not a primary aim of the various XLibris publications, nor are we attending to achieve free-form ink annotation.

3. The Digital Reading Desk

The existing literature has demonstrated the value of mark-up to support active reading (Adler et al., 1998). However, the actual use of annotation to support reading digital texts is low. This section describes the design and implementation of our improved digital reading and note-taking system, the "Digital Reading Desk," that extends our previous work on lightweight design (Pearson et al., 2009; Buchanan and Pearson, 2008). Although we are aware that reading on paper does not provide a panacea to reproduce to improve digital interaction, it does offer a proven, effective contrast (Landoni and Gibb, 2000). We anticipated that adopting the interaction style of paper as closely as possible would lead to a similar lightweight digital experience that would improve the effectiveness of electronic reading.

The Digital Reading Desk provides a rich user experience, with an interface for annotating documents with a minimum number of tools. Rather than use extensive menus and dialogs, we exploit direct manipulation and an extended workspace to minimise a user's interaction effort. We anticipate that minimising the cognitive attention being demanded by the tools will free more attention for the main note-taking task. The three main aims (next) of this design will be used later in the paper as criteria for evaluating the system's overall success.

3.1. Design aims

The main aims for the Digital Reading Desk are as follows:

A1: Provide an expanded workspace where users can position notes independent of the current open page (Section 1.2.1).

A2: Combine annotation and place-holding into a single tool (Section 1.2.2).

A3: Create drag-and-drop editing of notes to facilitate direct manipulation and encourage flow within the reading task (Section 1.2.3).

We will refer back to these aims throughout this article; they are numbered for easy cross-referencing. The implementation of these aims will be evaluated in the user study to test their contribution to the overall usability of the design.

3.2. Virtual desk area

As stated in Section 1.1, previous studies of annotation concluded that the space surrounding the document, whether large margins or desk space, empowers a user's note-taking. Although marginal space is seen in some software (e.g., Adobe Illustrator provides extra space surrounding its canvas area), it is seldom incorporated into document reading software. In order to improve the usability of such software we have included a virtual desk area as a backdrop for the document, which supplies extra space for the reader's notes.

As we are drawing on the behaviours of a physical book, it seemed consistent to present the book in a similar manner to paper. The PDF document being read is thus displayed in the same way as a physical book, as a double page spread is shown when the book is open, and a single page (either the front or back cover) when it is closed (Fig. 4). We anticipated that this consistency would cue users to interact with it like a physical book, and consequently to make better and more frequent use of the tools provided. It should be noted that the desk area belongs to the document itself; therefore opening a new document will give the user a new desk. This feature fulfils aim A1.

3.3. Unified Post-it tool

Paper is multi-functional: it can be used for notes or for place-holding. To reproduce this, we amalgamate these roles into a single tool. We follow previous work on visual placeholders (Buchanan and Pearson, 2008), where they protrude from the side of documents, like tabs in real books. Users can create notes and bookmarks from the *same* object, and removed any constraints on *where* the "Post-its" can be positioned, implicitly providing three options:

1. Completely within a document page.
2. Protruding from the document, causing it to act as a bookmark.
3. On the desk next to the document, or if the book is closed, Post-its can be placed in the area behind the document if desired.
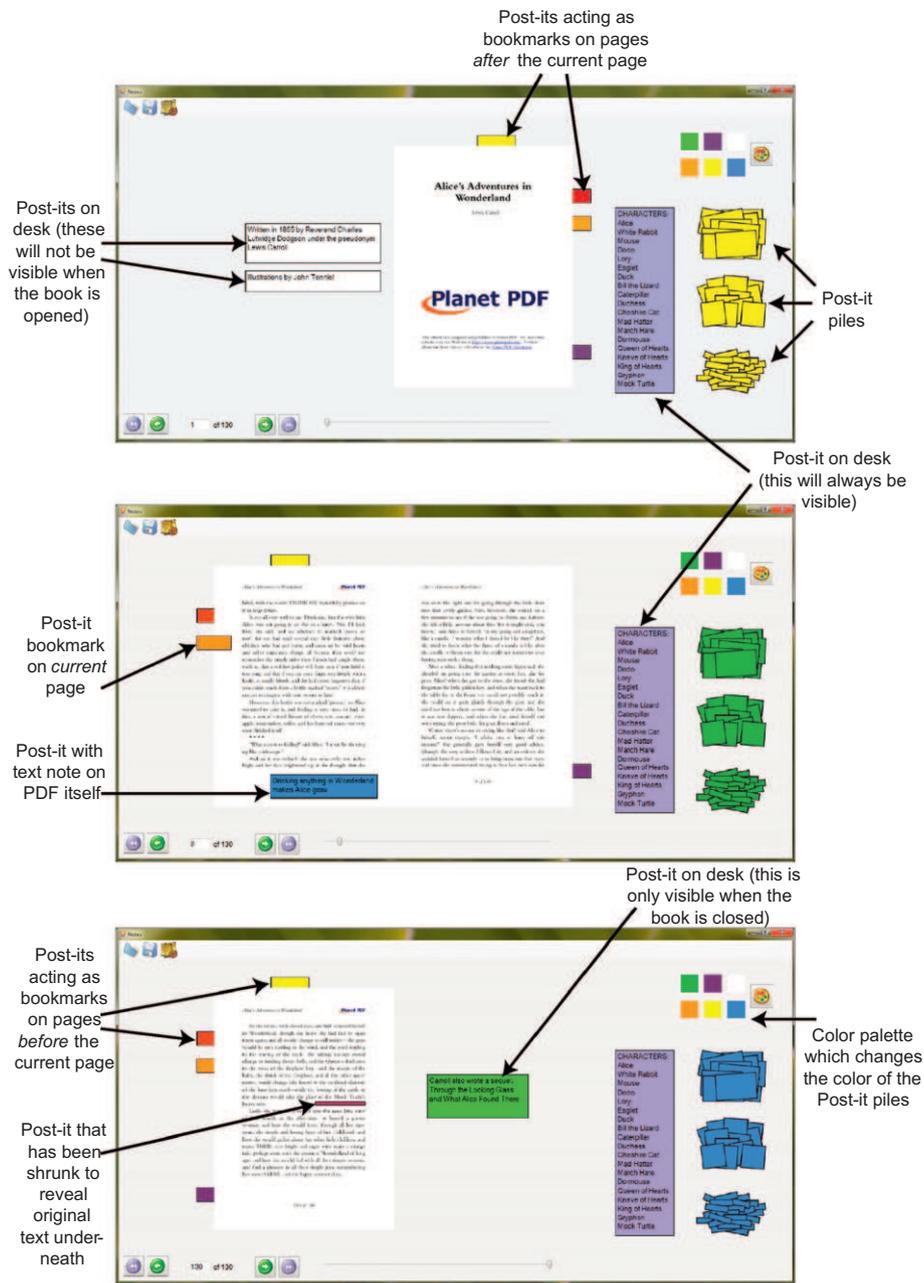
Fig. 4. Screen shots from the Digital Reading Desk.

Thus, one tool now performs three separate functions: making notes on specific pages (Point 1); making notes about the book as a whole (Point 3) and creating notes that also act as placeholders (Point 2). Post-its that also act as bookmarks not only navigate to the correct page when clicked, but also "flip" from one side of the book to the other depending on the current page. Post-its that bookmark pages that are *before* the current page appear on the left of the book, whereas those that are on pages *after* the current page appear on the right (much like the system shown in Fig. 3). See Fig. 4 for screen shots of the system. This is a potential solution to the problem of multiple, complicated tools (see Section 1.2.2) and fulfils Aim A2.

3.4. Drag and drop

To reduce on-screen menu clutter, we use a drag-and-drop style interaction for the creation and deletion of Post-its. On the right of the desk are three inexhaustible piles that can change colour using the palette at the top. To create a Post-it, the user drags one from the pile onto the document (Fig. 4). Notes are removed by dragging them back onto the pile. As well as the addition of text to the Post-its, they can be moved, resized or "lifted up" (to reveal text underneath), borrowing from and extending the behaviour of physical notes. All interactions are performed without menus: to add text is a double click, "lifting up" is a single click, and so on. This interaction addresses Aim A3, using direct manipulation and reduces the cognitive attention required by the tools, allowing the user to focus on the primary active reading task.

4. User study

This section reports a user study to assess the effectiveness of the Digital Reading Desk. The study compares our interface against two other designs that mirror common interactions of current digital document reader software. A small pilot study was conducted

before the one described below to refine both the system and study design.

4.1. Comparison systems

To provide appropriate systems to compare our design against, we implemented two additional interfaces. All three systems used the same basic features (e.g., display and page navigation) to provide consistent interaction, eliminate technical variables (e.g., rendering quality and speed) and avoid potential product bias etc.

The first baseline system, which we called traditional PDF (Fig. 6), is modelled on conventional PDF readers (e.g., Adobe Acrobat). Post-its can only be used as notes (i.e., they do *not* support bookmarking). They are created and deleted using pane dialogues: e.g., to create a note the user first selects its size and colour from options on the right of the window then click the New Note button to add it to the PDF (Fig. 5b). To delete a note the user must right-click over the note and select Delete . There is no usable desk area around the document, and notes cannot be placed outside the PDF "page." Once notes have been added, they provide the same basic functions as those on the Reading Desk (i.e., they can be moved, resized etc). Even this simple interaction is less demanding than most current reader software, where these features have to be accessed through menu navigation.

In what might be called the traditional PDF system, bookmarks are completely separate to notes and are stored in the same ordered tree used by most digital reader software (Fig. 5a). To add a bookmark, the user must navigate to the page then click New Bookmark . To delete or rename, they must use the right mouse button to access a short local menu. Note that these bookmarks are entirely user-created and do not include pre-defined bookmarks generated by the author or publisher (e.g., Chapter headings).

The second control system, called the traditional desk (Fig. 7), provides an intermediate point between the traditional PDF and reading desk systems. The traditional desk system is the traditional PDF system with the addition of a desk area. Notes and bookmarks

are again separate, but the notes can now be added to the desk as well as to the PDF itself.

4.2. Research questions

After implementing the Reading Desk interface, we laid down a set of research questions we hoped to answer via our comparative evaluation.

RQ1: Do the tools and interface of the reading desk better support users' note-taking than either the traditional PDF and traditional desk systems?

RQ2: Is drag-and-drop bookmarking (Reading Desk) easier to use than traditional, tree-listing of bookmarks (Traditional PDF and Desk systems)?

RQ3: Will incorporating both notes and bookmarks into one unified tool (Reading Desk) make the interface easier to use than two separate ones (Traditional PDF and Desk systems)?

RQ4: Will allowing users to place notes on a desk make it easier to make notes about the book as a whole (traditional desk and reading desk interfaces versus the traditional PDF presentation)?

RQ5: Will a visual approach to creating and deleting notes (Reading Desk) will improve their ease of use?

In order to answer these questions, our focus is on user ratings and remarks rather than using performance measures such as time to complete a task. This is because, as noted in Section 1.2, users of electronic reading software tend to do less note taking and express poor satisfaction with the interaction for creating and editing notes.

4.3. Procedure

We recruited 16 participants of mixed age/gender and high skill level (post-graduate degrees or above) to take part in the study. The study was undertaken in a quiet, well-lit room with plentiful natural light. Each session took an average of one hour. After an overview of the experimental procedure and ethical issues, the users were familiarised with the systems before completing 3 sets (one for each of the three systems) of 9 closed tasks (T1–T9)—see Section 4.3.1. To mitigate possible learning effects, the order in which the systems were presented to the participants, and the order of the task sets, was systematically rotated. The users were then asked to perform a different open-ended reading task with each interface. The equipment used consisted of a mouse and keyboard, plus a high definition (1920×1080 pixels) 17 inch (43 cm) widescreen display.

After the studies, each participant completed a short questionnaire that probed their current note-taking use. We then collected 15 subjective responses from our participants: five overall questions (Q1–Q5) and questions regarding the nine closed tasks (T1–T9). Each rating used a 5 point Likert scale from strongly disagree (1) to strongly agree (5). Finally, we conducted a semi-structured interview which questioned the participants on their experience of the systems and their impression of other note-taking techniques on paper and digital media. No video or audio recordings were taken during the study, though notes were taken by the observer. Screen capture was not possible, as this proved incompatible with the PDF rendering software. The final annotated files were anonymously stored for later use. The participants were given a $5 gift voucher in return for their time.

4.3.1. Study tasks

The study comprised two major parts for each interface: a set of ten closed tasks, and one large open task. The closed tasks were
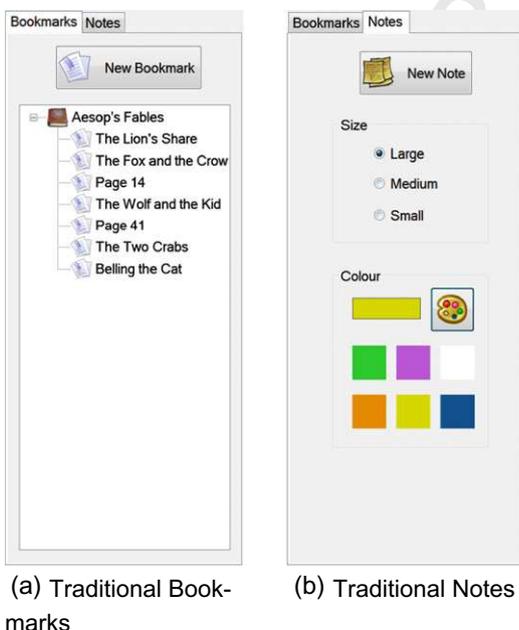
Fig. 5. The bookmark and notes interfaces on the traditional PDF and traditional desk systems.

(a) Traditional Bookmarks

(b) Traditional Notes
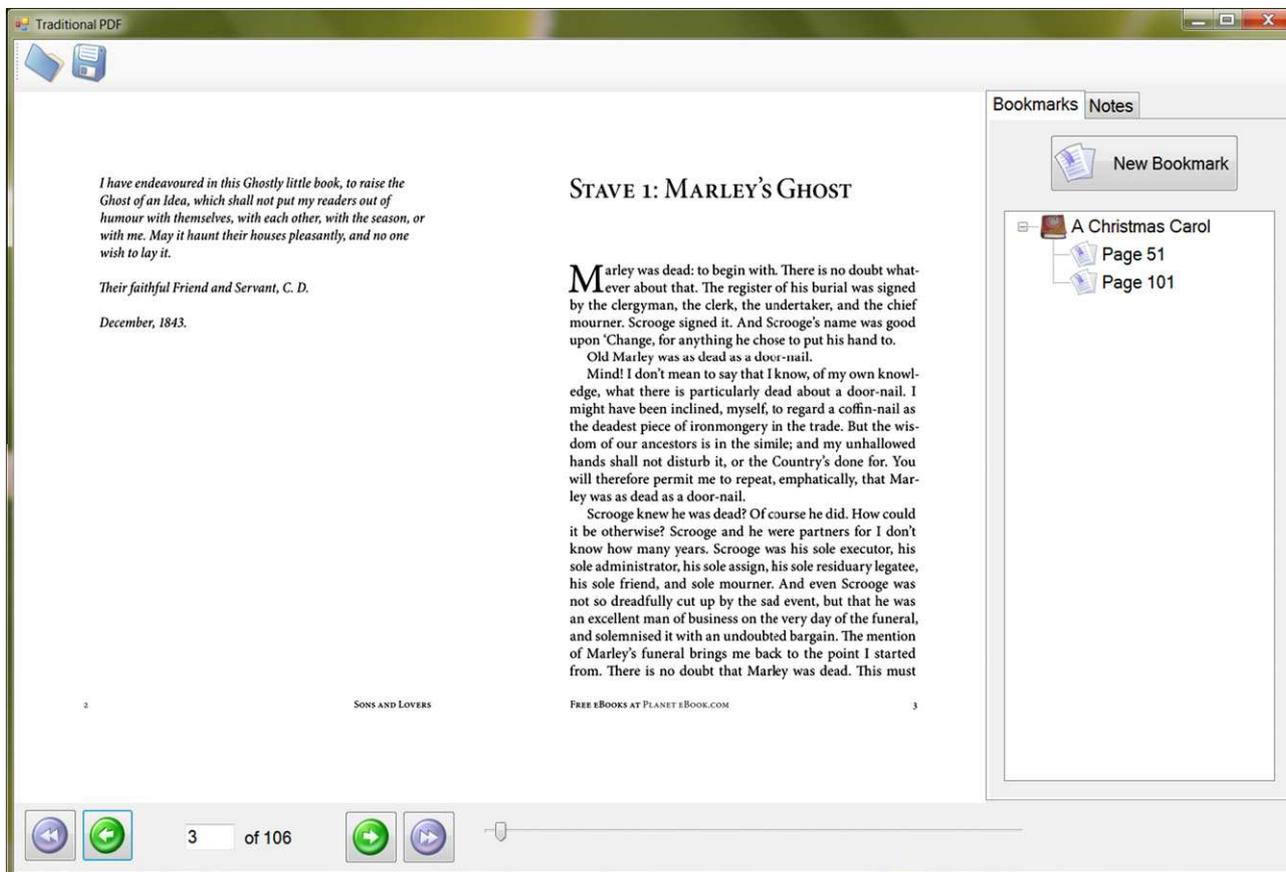
Fig. 6. The traditional PDF system.
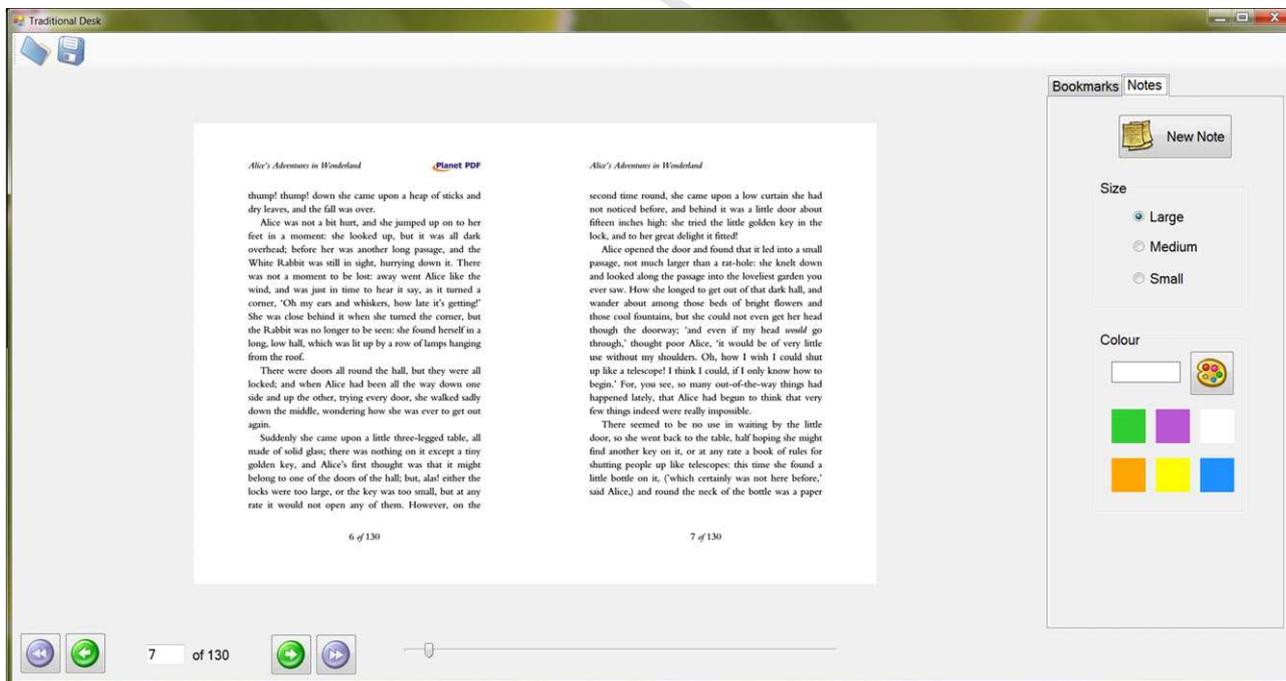
Fig. 7. The traditional desk system.

designed as learning aids and covered all the tools of each system, permitting a detailed measurement of specific low-level interactions. Each system experiment in the closed tasks used a different PDF document. The documents were assigned to systems in rotation to reduce any possible bias. The three documents used in the closed tasks used the same font type and size, and were around

108 pages each. The ten closed tasks varied slightly depending upon the exact PDF used, but used the same basic functions. These tasks correspond to the questions T1–T9 shown in Table 1. A typical closed task for a system within the study was:

> T4: Create a note for page 6 that says ''The name 'Scrooge' has entered the English vocabulary as a synonym for a miser.''

Due to their closed nature, this set of fixed tasks would not give us an accurate representation of how users would take notes in everyday life.

To obtain better information on "natural" user behaviour, the second part of a user's session with each interface was a larger open task. The open task encouraged the participant to mark up the documents in whatever manner they found most appropriate. We anticipated that any "genuine" difference in the interaction styles would result in distinct behaviours and output artefacts. The open tasks used three separate PDFs, again rotated between the interfaces for balance. The documents were of the same length (8 pages) and the tasks had a common structure. A sample task from a document entitled "Biofuels: Implications for food and agriculture" consisted of:

> Read the article carefully, finding strong points for or against bio-fuels, as well as any open questions that occur to you or any significant statistical data. Mark up the article with notes to help you explain your interpretation of the text to another person. Summarise the article with a brief paragraph that would explain your impression of it.

It is important to note that in an ideal world, we would undertake a longitudinal study of the Digital Reading Desk in use by appropriate users (e.g., academic researchers). It would be premature to embark on such a study before proving the system's potential benefit in a more controlled environment, and it is unlikely that users partaking in a longitudinal study would continue with a tool with key deficits, hence an initial laboratory study was a first step.

4.3.2. Study documents

The three PDFs used in the open tasks were designed and created by us. The topic of these three documents were varied within non-computer related topics and included multiple opinions and issues. The titles of the three PDFs used in the open tasks were: *Biofuels: Implications for food and agriculture*, *Video games and their effect on childhood obesity*, and *Tanning Beds: is bronzed skin worth the risks?* The reason for the short document length (8 pages) in the open tasks was due to the type of task being performed. Participants were asked to read each document carefully and make notes—longer documents might prove counter-productive by discouraging the main task of mark-up due to fatigue.

4.4. Study metrics

One major challenge we faced was that it is often complex to precisely assess the benefits of a new interactive paradigm. Ultimately, our hope was that our new interaction would encourage a change in the pattern of user behaviour, underpinned by lower mental effort, rather than make substantial time savings. Both our previous research, and our pilot study, demonstrated that time savings, whilst possible, were relatively small and, more importantly, unlikely to significantly change the patterns of user behaviour that have been seen as shortcomings of digital reading: that is, low rates of use, rather than problems centred on speed. Indeed, if we were successful in increasing digital annotation, we may actually see that the time taken to produce notes will *increase* due to an increasing number of annotations. If the prior research correctly ascribes higher levels of annotation with more attentive reading, a more successful design will encourage readers to spend more time writing their notes.

Mental effort is difficult to assess objectively, so we hoped ultimately to identify differences through the artefacts of open tasks. By studying *how* and *where* our participants made use of notes as well as the number of notes made, we are able to assess the usefulness of the different aspects of the system, such as the desk area. Using artefacts as an evaluative method is problematic however. In our study, would a higher level of note-taking activity reflect a superior interface? Alternatively would more note-taking in fact indicate that reading is being interrupted and users are calling more on notes to support a reading task that has been made more complex? In order to obtain a principled model for our study, we have drawn strongly upon Shipman's analysis of VKB (Shipman et al., 2004) who faced similar problems: the correct interpretation of artefacts and metrics for the task was imperfectly understood.

4.5. Results

The following section reports findings from the comparison study described above. Where appropriate, we have provided qualitative data from the participants to substantiate our claims. The quotes given below are a sample of the more relevant pieces of dialogue taken from the post-study questionnaires and the post-study interviews. The participant number of the user who gave the data is displayed in brackets after each quote (e.g. [PX]). The qualitative data was categorised and organised into simple codes by induction, organised by the relevant part of the semi-structured interview,

Table 1
Average subjective ratings (5 point Likert). X indicates a non-significant result.

	Traditional PDF		Traditional desk		Reading desk		ANOVA p
	Avg	SD	Avg	SD	Avg	SD	(F)
Q1: The interface was easy to use	3.56	0.89	3.94	0.57	4.75	0.45	<0.001 (13.34)
Q2: The tools were easy to use	3.5	0.89	3.81	0.75	4.56	0.63	<0.001 (8.14)
Q3: The tools were easy to *learn* to use	4.25	0.77	4.25	0.77	4.81	0.40	0.032 (3.72)
Q4: I would use this system out of choice	2.5	1.03	2.81	1.05	4.06	1.18	<0.001 (9.22)
Q5: This system mimics paper well	2.44	1.03	3.06	0.93	4.31	0.79	<0.0001 (17.12)
T1: It was easy to create new bookmarks	4.00	0.97	3.94	0.93	4.63	0.62	0.05 (3.19)
T2: It was easy to create new notes for a *specific page*	3.88	1.15	4.06	0.99	4.75	0.45	0.024 (4.05)
T3: It was easy to look up old bookmarks	4.31	0.95	4.31	0.95	4.31	0.60	**X** (0)
T4: It was easy to create new notes for the *book as a whole*	2.81	1.42	4.38	0.89	4.63	0.62	<0.001 (14.49)
T5: It was easy to look up old notes	2.63	0.96	2.63	0.96	3.63	1.26	<0.014 (4.68)
T6: It was easy to collapse and expand notes	4.00	1.09	4.00	1.09	4.00	1.09	**X** (0)
T7: It was easy to amend old notes	3.94	1.06	3.94	1.06	3.94	1.06	**X** (0)
T8: It was easy to delete notes	4.25	0.93	4.31	0.87	4.63	0.80	**X** (0.85)
T9: It was easy to delete bookmarks	4.19	0.98	4.25	0.93	4.56	0.63	**X** (0.87)

which followed the same overall structure as our quantitative feedback. We have primarily selected those responses that included comparison or rating utterances.

4.5.1. Subjective feedback

Table 1 reports the participants average subjective ratings. To assess for statistical significance we performed an ANOVA test on the 5 point Likert values.

General questions (Q1–Q5): The first general question posed to our participants was how easy to use they found the interface. The results of Q1 produced a highly significant preference for the Reading Desk interface proving that it is the easiest to use of the three. Similarly, Q2 (The tools were easy to use) and Q3 (The tools were easy to **learn** to use) both produced significant ANOVA results, also concluding that the Reading Desk's individual tool set is also superior to those in the alternative systems. These results indicate then, that RQ1 was satisfied with typical comments, such as the following, being consistent with our quantitative findings:

> ''The reading desk system was so intuitive, it acts as real world experience would lead you to expect.'' [P4]
>
> ''It fits in with the book style and look—I think even my mother could use this.'' [P3]

We also wanted to determine if participants felt the Reading Desk could be put to everyday use. Q4 (I would use this system out of choice) produced promising Likert values which proved statistically significant. This result is strengthened by user comments such as:

> ''The last one [digital reading desk] is pretty awesome, it is a big improvement over Preview for the Mac.'' [P9]
>
> ''It [the Digital Reading Desk] is cool—it makes logical sense, like real books.'' [P5]

The last of the general questions addressed how well our Reading Desk mimics the behaviour of paper. Although imitating the properties of paper may not be a perfect solution to the problems faced with digital mark-up tools, we feel it may enhance usability in some areas. Q5 then (This system mimics paper well), yielded highly significant Likert values which favour the Digital Reading Desk.

Specific tasks (T1–T9): As well as the five overall questions (Q1–Q5), the users were also probed about their thoughts on the closed tasks they performed (T1–T9).

The first task, T1 (It was easy to create new bookmarks) confirmed that users found creating new bookmarks easier with the Reading Desk's drag-and-drop method rather than the tree display. Participants' comments from the post-study interviews strengthened this finding:

> ''I prefer being able to drag and drop—it relates more to real life.'' [P9]
>
> ''It's more intuitive, you don't need to teach it because I already know how.'' [P14]
>
> ''Dragging is much better, I'd rather drag than click—it's far more fun.'' [P5]

During the pilot studies, it became clear that users tend to make two types of note: those that specific to a page (e.g., a note about a paragraph), and those for the document as a whole (e.g., a book summary). Where users place notes typically indicates of their type: for instance, notes about specific pages are usually placed on the page they relate to, whereas notes about the document as a whole are often placed either on the front cover, or on the desk to enable persistent access to the material. As a result, the tasks probing new note creation has been divided into two distinct questions: T2 (It was easy to create new notes for a *specific page*) and T4 (It was easy to create new notes for the *book as a whole*). We anticipated that Q4 would elicit useful feedback on the Desk element of our design; the observed behaviours and artefact placements are discussed later. Subjective user ratings for the creation of notes for a *specific page* produced an advantage for the Reading Desk, despite high Likert scores for all systems, and the ANOVA test was statistically significant.

For creating new notes for a *book as a whole*, ANOVA (followed by Tukey's HSD) shows a clear distinction between the two desk implementations (the Reading and Traditional Desk systems) compared to the Traditional PDF. This strongly suggests that the desk feature provides a useful area for placing notes about a document as a whole. These results, plus the data from participants' summary placement (Table 2) confirms RQ4, that placing notes on a desk will be easier for making notes about a whole book.

T2 and T4 explored the methods for Post-It creation. The three systems exploit two separate interactions: the Traditional systems (seen in Fig. 5b), use a click-button interface while the Reading Desk employs drag-and-drop. Both tasks T2 and T4 show differences between these two interactions, supported by user responses which included:

> ''Drag-and-drop was much easier to create new notes in different colours and sizes . . .you could see what it would look like before you place it.'' [P12]
>
> ''Dragging and dropping notes is more intuitive, why have menus when a picture [piles of Post-its] is better?'' [P3]

Despite the different interactions for bookmarks, task T3 (It was easy to look up old bookmarks) resulted in identical ratings for all three interfaces, a non-significant result. We conclude that there is no significant performance problem with the look up feature of our new bookmark interaction as the average Likert rating for all three systems was 4.31 out of 5.

In the two benchmark systems (Traditional PDF and Desk) the bookmark and note features were split into two distinct tools, while in the Reading Desk interaction they were merged. It thus seemed appropriate to ask, how easy users found it to look up old notes (T5: It was easy to look up old notes). The results produced a significant result favouring the Digital Reading Desk. Although several users suggested that an improved note look-up system (e.g., an additional togglable list) would prove beneficial, these results illustrate the popularity of the drag-and-drop interaction (Reading Desk) compared to the two traditional implementations.

RQ3 then, has also been confirmed using the subjective opinions of the participants including commments such as:

> ''It [the unified tool] was smooth, it simplifies the interface without sacrificing functionality.'' [P2]
>
> ''It's easier to act instinctively with the one tool as you are always reaching for the same things.'' [P2]
>
> ''It [the unified tool] makes life easier. In real life you wouldn't have a pile of Post-its and a pile of bookmarks you would just use a Post-it for both.'' [P14]

Table 2
Summary placement (% of participants).

	Trad PDF	Trad desk	Reading desk
Front Page	75%	19%	31%
Back Page	16%	3%	3%
Page Two	6%	0%	0%
Other	3%	0%	0%
Desk (always visible)	N/A	69%	57%
Desk (behind first page)	N/A	9%	6%
Desk (behind last page)	N/A	0%	3%
Desk (Combined)	N/A	78%	66%
Book (Combined)	100%	22%	34%

Tasks T6 and T7 produced the same mean scores for each system, and thus had non-significant results. These were, however, tasks where functionality was the same across all three systems, and similar results are both to be expected, and confirm participant neutrality. The final tests, T8 (It was easy to delete notes) and T9 (It was easy to delete bookmarks) yielded non-significant ANOVA results.

Although tasks T8 and T9 did not produce significant results, there is other evidence that the drag-and-drop creation and deletion of Post-its is more popular than the traditional "menu-and-click" method. While participants found both methods easy to use, they had a preference, as seen in Fig. 8. These results show that the majority (81%) of the participants felt that the drag-and-drop method was much better than traditional methods, (i.e., right click/delete). This is supported by the result of a χ^2 test with Yates' corrections: $\chi^2 (4, N = 16) = 35.875 \ p < 0.0001$. Participants' comments included:

"The drag-and-drop system was much the same as paper—chucking a note back on the pile is like screwing it up and throwing it away." [P3]
"Dragging notes back onto the pile is much simpler and more intuitive than fiddling with menus as this breaks me out of my train of thought." [P2]

Thus, RQ5 is partially supported by our results.

4.5.2. Patterns of use and user behaviour

As well as the post-study interviews and questionnaires, we also evaluated our participants' use of the three systems through observations of their interaction in the open tasks, and analysis of the final marked-up documents. From this data, we identified a range of interesting mark-up behaviours.

The use of bookmarks differed significantly between the systems: 75% of participants used them on the Reading Desk system; while only 25% and 19% used them on the Traditional PDF and Desk systems respectively. This threefold increase in bookmark use strongly suggests that the unified Post-it tool actually encourages their use. The results of a χ^2 test on this data gave significant results: $\chi^2(2, N = 16) = 7.684 \ p = 0.021$.

Secondly, the uptake of the desk area was promising, with many participants making use of it when available. The behavioural data from users placement of document summaries is one example of this. Table 2 reports the results from task T4 (making notes on the *book as a whole*) as well as summarising behaviour in the open tasks. To assess statistical significance, we performed a χ^2 test with Yates' corrections (primarily due to the low frequency of "back page" option). Not surprisingly, the difference between the Reading Desk and Traditional Desk systems is not-significant, however the results of these systems versus the PDF system produced a highly significant result of $\chi^2(4, N = 16) = 26.93 \ p < 0.0001$. It is clear from data that the desk is the preferred position for document
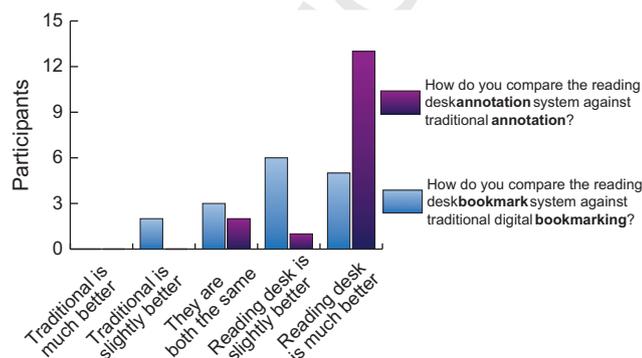
summaries. There is further support from the post-study interviews:

"I always make summaries separately so being able to put them on the desk is nice as you can always see it and make notes as you go." [P8]
"The desk is good because it is always in arms reach." [P11]

Some circumstances in which a desk area was said to be useful included: storing summaries of whole books rather than specific pages; taking notes to support direct reading (e.g., "read this first," or a task reminder); for common definitions; character biographies.

In terms of RQ2, Fig. 8 reports preferences for bookmark presentation: traditional (tree-list) versus drag-and-drop (unified Post-it tool). While 69% of the participants rated the unified Post-it tool to be slightly or much better than the traditional tree list, this result was not statistically significant ($p = 0.12$). Our results are thus supportive rather than conclusive. While tab style bookmarks are, overall, more popular, participants reported useful properties of the tree bookmark list, particularly when looking up notes. One commented:

"Creating and reading should be drag-and-drop; and browsing should be a list." [P14]

—which indicates that the overview of a tree-list is more beneficial for reviewing notes than the visual tab interface, while the reading desk interface is better for creating them. A likely improvement might be incorporating a togglable list into the reading desk interface.

4.5.3. Results summary

The results of our study have demonstrated a number of advantages created by the use of a larger workspace around the document (the desk) and the provision of drag-and-drop control of notes and bookmarks in a unified Post-it tool. This general picture is underlined by the participants' response to Q5, which yielded a clear outcome. Even more importantly, user behaviour changed in the Reading Desk design: specifically, users made more use of notes and bookmarks. This change in behaviour has proved an elusive goal for researchers in the past. Our current design can likely be improved, and the causes of success more closely examined.

5. Discussion

The data from our study has provided evidence that there are a number of impacts from the adoption of a direct manipulation interaction style. There were differences in user behaviour during the open tasks. The key advantages of the system include: the desk area surrounding the document, the unified Post-it tool (including tabbed bookmarks) and the simplified drag-and-drop create and delete techniques.

The role of margin space within digital document readers has been addressed in a previous study (Pearson et al., 2009). We used that work as a basis for introducing the desk area. Observed behaviour in the study, and analysis of the subjective feedback, confirmed the utility of the desk area, and proved RQ4.

We have also identified several key advantages of the unified tool, including its ability to support bookmarking and notes in one tool. The unified system was also rated easier to learn. Participant behaviour and feedback both supported and confirmed RQ3.

The participants also expressed their preference for the bookmark 'tab' system, an enhancement of an earlier interface (Buchanan and Pearson, 2008) (Fig. 3). As with the earlier design, users can track pages of interest by means of coloured tabs that stick out of the sides of the document. Unlike the earlier system, which used page order to determine the visual position of a bookmark (i.e., bookmarks at the top are on lower page numbers than those

Fig. 8. Subjective reviews of the annotation and bookmark portions of the systems.

on the bottom), the drag-and-drop system, by allowing users the freedom to drag notes anywhere, loses this sense of order. This may be a contributory factor for the participants view that a hybrid bookmark interface with both drag-and-drop tab and togglable tree-list features would be beneficial.

5.1. Lessons learnt

The challenge of improving the interaction of document reader software is not trivial. Progress has been slow despite years of concerted effort (Marshall and Bly, 2005; O'Hara and Sellen, 1997; Schilit et al., 1998b). This research aimed to improve the use of annotation and place-holding features in a digital document reader, by using one dual-purpose tool where currently there are two separate tools. The unification of place-holding and annotation mirrors the practices of working with physical documents (Marshall, 1997; Adler et al., 1998). We hypothesised that providing a single tool would produce user behaviour that more closely reflects the behaviours seen on paper. This unified approach extends previous work (Buchanan and Pearson, 2008; Pearson et al., 2009; Schilit et al., 1998b), and applied the general lessons learnt, providing the free placement of notes and an extended workspace beyond the nominal paper bounds of the document display.

Our evaluation demonstrated several advantages to our design. We did not expect significant differences in the time performance of adding notes in the various interfaces, nor did we anticipate that a drag-and-drop interaction would be dramatically easier to use than a menu based method. Rather, we anticipated that the different interaction styles would result in a greater use of notes and placeholders in the unified tool, particularly in open tasks. This difference was observed in our study. The evidence gathered is not, we believe, conclusive proof that a single tool for annotation and place-holding is the optimal design, but there is a clear improvement over having two separate tools. It is our conjecture that a single tool will become more familiar to the user and through this will require less cognitive attention. It will require further experimentation to prove this conclusively.

A continuing goal for our research is to understand how to make interacting with digital documents lightweight. We do not believe that all interaction with paper is lightweight: obviously, searching for a word deep in a document is more laborious than when done on a computer. One understanding of lightweight, then, is that secondary, supportive tasks such as annotation become lightweight when they require minimal cognitive attention. Direct manipulation may support a reduction in cognitive load when compared to current designs. However, lightweight is more than a single issue. Unified tools may also assist in the reduction of cognitive effort; so may the use of colour and other visual cues. Our current designs thus contribute to a better understanding of how to make digital interaction with documents lightweight.

5.2. Reader appliances

As digital documents and reading appliances become more common, the need for efficient reader software becomes a necessity (Pearson et al., 2010). A fundamental problem with specialist reading devices is their lack of screen real-estate. Where the limited display is dedicated to the document, this provides minimal space for notes and placeholders. In addition, the limited button controls of reader hardware makes it difficult to deliver a drag-and-drop interface.

There have been several attempts to improve the usability of eInk hardware. Chen et al. (2008) discuss the benefits of a dual-display e-book reader. The authors describe various problems with current methods, including differing affordances between paper and digital, and the lack of screen space. To remedy these prob-

lems, they introduced a second display and support embodied interactions such as folding, flipping and fanning that they describe as lightweight navigation. Their user study proved that the lightweight aspects of the design supported reading, but the embodied interactions did not.

Returning to the Digital Reading Desk, our current implementation is set in the context of a desktop environment. Some of the design decisions we have made rely on the ready availability of larger expanses of screen estate. However, on many reading appliances, this is in shorter supply. The principles that underpin our design could readily inform superior interactions for reading appliances: for example, note tabs can be rendered in a very space-efficient form, retaining many of their advantages without significantly reducing the available reading space. Not all parts of our design will translate so readily. However, the underlying principles may be used to inform different designs that would be more effective where screen estate is more limited.

5.3. Annotation types

In this paper, we have specifically focussed on written annotation in the form of Post-it notes, and on place-holding. There is are, clearly, other forms of mark-up which we did not explore (e.g., highlighting). Some of our key principles—for example, using one tool for multiple purposes—could also be applied to other annotation tools. In the case of a highlighter, abstracting to a pen-like tool would permit both drawing and highlighting in one. For the purposes of our immediate research, we limited ourselves to a smaller set of mark-up methods to avoid casting our research too widely to draw meaningful conclusions. It is now, clearly, an avenue of future work to expand the range of tools whilst re-testing and re-applying the principles this first stage of research has proven. For our existing methods, and future tools, one key advantage of digital texts remains—that the "original" does not have to be permanently defaced by the user's annotation.

We have previously demonstrated the value of temporary placeholders (Buchanan and Pearson, 2008), and these are of particular value when a user is switching between a number of documents. In the digital realm, the potential of supporting interlinking comments between documents has yet to be fully explored, and there are few established methods on paper that have proved at all common in use. This is one area where our general lightweight approach may prove useful in creating novel interactions that do not simply mimic a paper precedent. The desk area is one potential feature of our design that could be used across and between documents, for example.

5.4. Possible future work

The information gathered during the post-study interview has given leads for improving the current design. For example, several users commented that the two-dimensional display of the text—particularly the edge of the book—made it hard to see where the Post-it bookmarks were positioned within the document (i.e., which ones were close to the front/back). To remedy this shortcoming, a visual cue could give a better indication of where each bookmark is placed in relation to others in the document. In the physical world, a number of subtle cues combine to give such information—like the changed hues of Post-its with changing depth, and touch. However, literal reproduction of this in digital form may be either ineffective or impossible. Potential digital solutions to this issue include adding dots to the navigation bar, similar to the interface presented by Byrd (1999), or a graph visualisation as suggested by Harper et al. (2003).

It would be beneficial to conduct a more in-depth study of the system, when using longer documents and over a longer period

of time. This would obtain a more naturalistic view of how users make use of the tools in their everyday mark-up tasks. A diary study over several days or weeks, or a longitudinal log analysis could both contribute to such a study.

Finally, at present, the desk is connected with the PDF that is open, meaning that opening a new PDF will give you a fresh desk. Although this is useful, it may also be worth considering the need for a persistent "meta-desk" that would mimic physical practices more closely, as some tasks may require the need for notes to be kept across multiple documents (Marshall, 2008).

6. Conclusions

It is well established that digital document mark-up tools are used less than their paper equivalents (O'Hara and Sellen, 1997; Marshall, 1997; Marshall and Brush, 2002). All note-taking tools can be seen as aids to a primary task, thinking of what to write. It is essential to the overall flow (Csikszentmihihalyi, 1990) that using these secondary tools is intuitive and tacit, to maximise the attention available for the main task. In the physical world this is usually the case. Scribbling notes on a Post-it, then placing it on a particular passage of text, is so straightforward that users often do it without thinking. In contrast, the equivalent digital tools are far less intuitive and suffer from low rates of use.

This article builds on previous research (Buchanan and Pearson, 2008; Pearson et al., 2009) to design digital mark-up tools that mimic the affordances of paper. We reduced the cognitive effort these tools demand, and increased the user's attention on the main reading task. Our improved design includes two key features: drag-and-drop Post-its that support both annotation and place-holding into one tool and a "desk" area that provides a static workspace around the document. We also reduced the menu system by exploiting drag-and-drop interactions for creating and deleting Post-its.

Our comparative study concluded that the interaction we created, using unified direct manipulation interaction combined with an extended workspace was preferred by participants. This supported our conjecture that a single unified tool for place-holding and note-taking would minimise mental effort. The feedback of participants also uncovered the value of the extended workspace for placing summaries and on-going notes to support longer term tasks. Most strikingly, we achieved an increase in the use of note-taking and place-holding tools in the preferred interface. A survey of the literature failed to uncover a similar outcome in any other study. Though further work is required to continue to improve on these results, we believe that this is an important, if early, step towards an understanding of how to deliver lightweight (Marshall and Bly, 2005) tools for digital active reading.

References

Adler, M.J., Doren, C., 1972. How to Read a Book: The Art of Getting a Liberal Education. Simon and Schuster.

Adler, A., Gujar, A., Harrison, B.L., O'Hara, K., Sellen, A., 1998. A diary study of work-related reading: design implications for digital reading devices. In: CHI '98: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, pp. 241–248.

Buchanan, G., Pearson, J., 2008. Improving placeholders in digital documents. In: Research and Advanced Technology for Digital Libraries. Lecture Notes in Computer Science, vol. 5173. Springer, Berlin/Heidelberg, pp. 1–12.

Byrd, D., 1999. A scrollbar-based visualization for document navigation. In: DL '99: Proceedings of the Fourth ACM Conference on Digital libraries. ACM, New York, NY, USA, pp. 122–129.

Chen, N., Guimbretiere, F., Dixon, M., Lewis, C., Agrawala, M., 2008. Navigation techniques for dual-display e-book readers. In: CHI '08: Proceeding of the Twenty-Sixth Annual SIGCHI Conference on Human Factors in Computing Systems. ACM, New York, NY, USA, pp. 1779–1788.

Csikszentmihalyi, M., 1990. Flow: The Psychology of Optimal Experience. Harper and Row.

Czerwinski, M., Smith, G., Regan, T., Meyers, B., Robertson, G.G., Starkweather, G., 2003. Toward characterizing the productivity benefits of very large displays. In: Human–Computer Interaction INTERACT: IFIP TC13 International Conference on Human–Computer Interaction. INTERACT '03. IOS Press.

Dillon, A., 1992. Reading from paper versus screens: a critical review of the empirical literature. Ergonomics 35 (October), 1297–1326.

Dix, A., 2007. Designing for appropriation. In: BCS-HCI '07: Proceedings of the 21st British HCI Group Annual Conference on People and Computers. British Computer Society, Swinton, UK, pp. 27–30.

Frommholz, I., Fuhr, N., 2006. Probabilistic, object-oriented logics for annotation-based retrieval in digital libraries. In: JCDL '06: Proceedings of the 6th ACM/IEEE-CS Joint Conference on Digital Libraries. ACM, New York, NY, USA, pp. 55–64.

Harper, D.J., Koychev, I., Sun, Y., 2003. Query-based document skimming: a user-centred evaluation of relevance profiling. In: Advances in Information Retrieval. Lecture Notes in Computer Science, vol. 2633. Springer, Berlin/Heidelberg.

Landoni, M., Gibb, F., 2000. The role of visual rhetoric in the design and production of electronic books: the visual book. The Electronic Library 18 (3), 190–201.

Liesaputra, V., Ian, W.H., 2008. Seeking information in realistic books: a user study. In: JCDL '08: Proceedings of the 8th ACM/IEEE-CS Joint Conference on Digital Libraries. ACM, New York, NY, USA, pp. 29–38.

Liesaputra, V., Witten, I.H., Bainbridge, D., 2009. Searching in a book. In: Research and Advanced Technology for Digital Libraries. Lecture Notes in Computer Science, vol. 5714. Springer, Berlin/Heidelberg, pp. 442–446.

Malone, T.W., 1983. How do people organize their desks? Implications for the design of office information systems. ACM Transactions on Information Systems 1 (January), 99–112.

Marshall, C.C., 1997. Annotation: from paper books to the digital library. In: DL '97: Proceedings of the Second ACM International Conference on Digital Libraries. ACM, New York, NY, USA, pp. 131–140.

Marshall, C.C., 2008. Collection-level analysis tools for books online. In: BooksOnline '08: Proceeding of the 2008 ACM Workshop on Research Advances in Large Digital Book Repositories. ACM, New York, NY, USA, pp. 41–44.

Marshall, C.C., Bly, S., 2005. Turning the page on navigation. In: JCDL '05: Proceedings of the 5th ACM/IEEE-CS Joint Conference on Digital Libraries. ACM, New York, NY, USA, pp. 225–234.

Marshall, C.C., Brush, A.B., 2002. From personal to shared annotations. In: CHI '02: CHI '02 Extended Abstracts on Human Factors in Computing Systems. ACM, New York, NY, USA, pp. 812–813.

Marshall, C.C., Price, M.N., Golovchinsky, G., Schilit, B.N., 1999. Introducing a digital library reading appliance into a reading group. In: DL '99: Proceedings of the Fourth ACM Conference on Digital Libraries. ACM, New York, NY, USA, pp. 77–84.

O'Hara, K., Sellen, A., 1997. A comparison of reading paper and on-line documents. In: CHI '97: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM, New York, NY, USA, pp. 335–342.

O'Hara, K., Smith, F., Newman, W., Sellen, A., 1998. Student readers' use of library documents: implications for library technologies. In: CHI '98: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, pp. 233–240.

Pearson, J., Buchanan, G., Thimbleby, H.W., 2009. Improving annotations in digital documents. In: ECDL. Lecture Notes in Computer Science, vol. 5714. Springer, Berlin/Heidelberg, pp. 429–432.

Pearson, J., Buchanan, G., Thimbleby, H., 2010. Hci design principles for ereaders. In: BooksOnline '10: Proceeding of the 2010 ACM Workshop on Research Advances in Large Digital Book Repositories. ACM, New York, NY, USA.

Pearson, J., Buchanan, G., Thimbleby, H., 2011. The reading desk: applying physical interactions to digital documents. In: CHI '11: Proceedings of the 29th International Conference on Human Factors in Computing Systems. ACM, New York, NY, USA, pp. 3199–3202.

Schilit, B.N., Golovchinsky, G., Price, M.N., 1998a. Beyond paper: supporting active reading with free form digital ink annotations. In: CHI '98: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, pp. 249–256.

Schilit, B.N., Price, M.N., Golovchinsky, G., 1998b. Digital library information appliances. In: DL '98: Proceedings of the third ACM Conference on Digital Libraries. ACM, New York, NY, USA, pp. 217–226.

Sellen, A.J., Harper, R.H., 2003. The Myth of the Paperless Office. MIT Press, Cambridge, MA, USA.

Shipman, F.M., Hsieh, H., Moore, J.M., Zacchi, A., 2004. Supporting personal collections across digital libraries in spatial hypertext. In: JCDL '04: Proceedings of the 4th ACM/IEEE-CS Joint Conference on Digital Libraries. ACM, New York, NY, USA, pp. 358–367.

Shneiderman, B., Plaisant, C., 2004. Designing the User Interface: Strategies for Effective Human–Computer Interaction, fourth ed. Pearson Addison-Wesley.

Wu, C.-S.A., Robinson, S.J., Mazalek, A., 2008. Turning a page on the digital annotation of physical books. In: TEI '08: Proceedings of the 2nd International Conference on Tangible and Embedded Interaction. ACM, New York, NY, USA, pp. 109–116.