

Real-Time Document Collaboration Using iPads

Jennifer Pearson¹, George Buchanan²

¹ Future Interaction Technology Laboratory
Swansea University
Swansea
csjen@swan.ac.uk

² Centre for HCI Design
City University
London
george.buchanan.1@city.ac.uk

ABSTRACT

This paper presents a real-time collaborative interface for reading groups. We exploit the light and portable properties of iPad hardware to facilitate an intimate interaction experience that gives users the ability to simultaneously mark up their own copy of a document, while instantly viewing notes made by other members of the group. We also introduce a tool that allows group members to quickly ‘point out’ information within a document to other members without physically gesturing with their hands. We finish with a short explanation of a study proposed to test the implemented system.

Categories and Subject Descriptors

H.5.2 [User Interfaces]: input devices and strategies; interaction styles; prototyping

General Terms

Design, Experimentation, Human Factors

Keywords

Annotation, Collaboration, Bookmarking, Documents

1. INTRODUCTION

Reading is not usually considered a social activity - being engrossed within your own copy of a novel does not promote collaboration. However, although it is unusual for several people to sit and read a book together, readers will often meet to talk about their experience and understanding of books that they have read. This is seen, for example, in study groups and book clubs.

Paper is the traditional and established medium for reading. When reading in a group environment, users will often make use of individual copies of the same document which, in itself, has clear advantages, e.g. everyone within the group

gets to see and make notes on their own copy of the document. These advantages are seen particularly with individual articles (perhaps downloaded and printed), but even with large books, the medium is readily manipulated and does little to impede the social interaction between people. The readers can sit in whatever space is available, without being constrained by the technology, and with minimal physical intrusion from the reading medium. While these benefits enhance the collaborative experience, however, they do give rise to some practical problems. For example, referencing a specific section can be tricky - the reader who is guiding their colleagues to a particular place must identify its physical location, and this can give rise to co-ordination problems. A second issue is that the readers’ individual notes are separate and sharing ideas can lead to considerable duplication of effort.

Using a single shared resource has limitations, whether using a computer or paper. Collaborative web search is a task closely related to collaborative reading, and research has shown that problems arise with both the human-computer interaction and the use of physical space [2] when using a single PC. Desktop computers restrict the physical space, and this limits the area available for co-ordination between users, and their collective control of the primary (search) task. Using a single resource is noted as encouraging a hierarchical group dynamic, with one user dominating.

Recently, computer technologies have multiplied, and there are large interactive displays that overcome some of the physical constraints of the relatively small display of a traditional desktop or laptop PC. This physical layout has also been examined in the context of collaborative search, covering interactive whiteboards, table-top PCs and similar devices [5, 9, 13]. These technologies can promote simultaneous collaboration, and permit everyone to see a shared display that co-ordinates the group’s view of the media being worked with - i.e. the referencing problem is by-passed through the use of a single view. However, this does not give each person an individual view.

This paper introduces a collaborative annotation system that uses a set of individual iPads. The size and shape of the hardware resembles clipboards and allows users the space to sit around the same desk while at the same time giving them their own document to work on. The wireless connectivity of the devices allows users to contribute easily (and not necessarily vocally) via real-time changes to all iPads in the group. This method not only facilitates simultaneous annotation, but also keeps users up to date on all notes made within the session. It supports both individual and group

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

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

reading, as naturally arises when readers have their own individual printed text. The following sections describe the background and design of such a system along with a short conclusion of the expected benefits.

1.1 Scenario

Kris, Joe, Sian and Emma are studying for an exam. They each have a copy of the course text on their individual iPads and are discussing the material around a table. Joe notices that there is a section within one of the chapters that is not fully explained. He read a web article on this topic earlier in the day and decides to make a note clarifying the information for the benefit of his friends. Immediately after he makes the note, it automatically appears on the others iPads along with a note at the bottom stating that Joe has added a note. Later in the session, the group are discussing a past exam question and decide to look up the answer within the document. Kris finds the appropriate page and paragraph and ‘points it out’ to the other three via the tools within the system. Joe, Sian and Emma’s iPads instantly recognize that Kris is pointing something out and offers them the chance to automatically jump to the right page.

2. BACKGROUND

Before we discuss the design for our system, it would be useful to take a look at similar collaborative systems to gain a better understanding of the techniques used for information sharing within working groups.

The work in [6] describes the importance of real-time collaboration tools for digital ink applications. It tests commercially available peer-to-peer ink-enabled software with pen-based interaction tools in a bid to identify the most efficient features. The hardware used in the study consisted of pen tablet monitors with separate CPU and peripherals and “hybrid” tablets with detachable keyboards: two devices that require a large amount of desk space and are significantly immobile. The main conclusion drawn from this study was the challenge and importance of facilitating deixis: a ubiquitous technique whereby an idea is demonstrated by speaking about it and pointing to it simultaneously.

This now leads us onto the problem of shared references [7]: i.e. if one person refers to something specific in a document how does everyone in the discussion know they are looking at the same thing? This can be achieved (somewhat clumsily) in paginated documents like books by pointing out the section with your finger, e.g. “go to page 22, paragraph 4”, but can prove complicated if people in the group are reading different editions or versions of the text. This method of referencing can also cause problems for electronic documents that lose their sense of location, such as eReaders where the page numbers alter depending upon the zoom level on the device, or documents in HTML format. Clearly then, using page numbers in this context is not a consistent method of referencing. It is for this reason, that the collaborative iPad system will make use of document locations as opposed to pages much like the format used by the Amazon Kindle software. (See Section 3.3 for details).

[11] investigates face-to-face or co-present group collaborations and formulated the model known as Single Display Groupware (SDG). This model is defined as a single user interface that allows simultaneous collaboration by multiple users via numerous input devices. There have been several attempts to improve the collaborative process by incor-

porating multiple input devices on the same display. [2], proposes a system to improve the often solitary task of web-search via a co-located collaboration group. The system known as *CoSearch* uses readily available devices such as mice and mobile phones to aid in traditional shared computer tasks. The main benefit of the system was that although the group members were gathered around the same computer, they each had facility to control it while at the same time maintaining group communication and awareness. Thus, all members of the group could a) see the same collaborative document and b) make changes to it. This is similar in part to what we are trying to achieve with our iPad system. The difference here is the interaction side: we want all users to feel they are working on their own copy of a document (hence the individual iPads) while at the same time allowing them real-time access to others notes.

Whilst much of the research on group collaboration focuses on table top displays and SDG, some uses more portable ‘paper-like’ devices such as tablet PCs. Hinckley et al [4] proposed a novel collaboration technique that connects together two tablet PCs fitted with physical sensors. The sensors are used to identify switches between shared and individual tasks, and to identify physical gestures that trigger document sharing.

There are also several commercially available systems that focus on collaborations on Apple portable hardware. Crowd-Map from Barking Minds¹, allows users to build mind maps collaboratively on their iPads or iPhones. The tool, which operates in real-time, has been designed for the organization and sharing of mind maps within working groups. Another system that makes use of ‘live’ collaboration using iPads and iPhones is letsannotate² that allows users to view and annotate PDF files. The portable iPad version is currently only in its beta stage, and does not run as an application, but through a web browser.

3. SYSTEM DESIGN

The aim of the collaborative iPad system is to mimic the way in which a group of users work together using paper based tools while at the same time giving the benefits of a single computerized machine. Specifically, we want to give users the freedom of annotating their own copies of a document, while at the same time allowing them instant access to notes and marks made by other members of the group. Just like in the scenario in Section 1, iPads take the place of physical paper copies of collaborative material.

3.1 Joining a Session

The design behind this system starts with every member of a working group connecting to a central server using their personal iPads and opening a copy of the collaborative document. Collaboration not only focuses on the exchange of information, but also the long-term persistence or storage of the information afterwards, e.g. oral information is not retained unless it is recorded. Therefore, to ensure each user is left with a marked up document at the end of the session, these documents are stored directly on each user’s machine.

To uniquely identify each user on the system, the iPad’s MAC address is used. For more personal identification in group sessions, users are requested to enter a ‘nickname’

¹Available at <http://crowdmap.barkingminds.com>

²Available at <http://letsannotate.com>




Figure 1: Collaborative iPad Design

which will then be used to distinguish themselves from others.

Users can come and go frequently from group collaborations. It is for this reason that the sessions act in the same way as typical chat rooms: i.e. users can log in and out. To ensure that late joiners contain the most up-to-date version of the collaborative document, the server records all changes made to it throughout the entire session and updates new users on the system. In addition to this, any annotations made to the document while a user is offline can be included into a working group by opting to do so at login time.

3.2 Annotations

When an annotation or note is made by a user in a group, their iPad sends all its information to the server, which then distributes the content back to some or all of the other members. The default distribution list can be set at the start of the user's session. Once an individual iPad receives an annotation from the server, it updates its copy of the collaborative document and informs the user of the change via the 'chat' box: e.g. "User3 has made a note on page 5". Although the document changes on a user's machine, the program will not automatically shift focus to the new annotation. It allows quick navigation to a new annotation through clicking on its entry in the 'chat' box. As with traditional chat clients, a record is kept of each new entry in a scrollable list until the session ends.

All details about an annotation are recorded by the system, including the time it was made, and by whom. This allows a variety of different interactions: e.g. a timeline of annotations can be viewed from the main menu, and selecting a current annotation can bring a pop-up of who created it. Additionally, the system can restrict who can delete annotations depending upon the preferences set by the group at the start of the session.

3.3 "Look At This" Queue

The other useful feature of the system is its ability to 'point out' things to users without actually marking the document. Often in a group collaboration session, users will want to indicate items and locations to other group members. For example, "look at this figure on page 45". When working on a single document or machine, this process is straightforward: the user can simply point out the section they are looking at with their finger. Assuming all members of the group have visible access to the shared document or screen, they will easily see the required information.

When collaboration occurs over multiple copies of the same document, however, this process becomes more troublesome, not least because there are now several copies to work with. Clearly then, simply pointing to the correct area on all copies would be tedious and time consuming and as we discussed in Section 2, attempting to describe a position in the text can prove difficult when dealing with different book editions or non-paginated documents.

To deal with this, our system has applied document locations that will correspond to specific points within the text, and will not alter on different zoom levels. We use these locations to implement a 'look at this' queue that allows users to 'point out' specific areas of the original document or annotations to other members of the group. These links to specific points in the text will not permanently mark the document and are stored in a queue on the right of the screen (see Figure 1). To avoid constant document page changing whenever another group member points something out, the system only navigates to a "look at this" entry when the user clicks on an item from the queue.

3.4 Architecture

The architecture of our iPad collaborative reading system is




Figure 2: Collaborative iPad Design

seen in Figure 2. There are a number of iPad devices (left and right of figure), connected together via a library server (center). When a user on a single device makes a gesture or annotation that is to be shared with other readers, this is forwarded to the library server as an event ①. The local PDF rendering code displays the annotation over the current PDF display ②. The library server then executes two steps (one of which is optional). First, the library server may record the mark (if it is a ‘permanent’ rather than ‘transient’ annotation) ③. In this case, the annotation will have a permanent ID set, ④ that will be later communicated back to the original iPad device ⑦.

Second, it must forward the message content to the other iPad devices in the group: to do this it communicates with the group service ⑤, receiving the IDs of the other iPads in the group ⑥. The service then sends the annotation to each group member ⑧, with the PDF renderers updating their display accordingly ⑨. Messages are sent in XML over TCP/IP sockets, and could be used with PDF rendering services on desktops, laptops and other devices.

Storing annotation records on the library server means that all the notes taken on a particular document can be recovered by later users. When reading a document a user can choose to view some or all of the annotations available (e.g. only their own notes, or public notes of colleagues). The diagram does not indicate the publishing controls used. This is a standard publish/subscribe model, but any alternative publish/subscribe model could be used instead with minimal re-engineering.

Prior work has reported database architectures for annotation (e.g. [1, 3]). Just as the publication controls used by our system can be adapted to work with alternatives, our live distribution of annotation events can be applied to various annotation database architectures. Our use of XML communication provides a neutral infrastructure upon which different annotation and PDF rendering services can be run.

3.5 Co-Located versus Remote

Our system has been designed as a collaboration device for users sitting in close proximity to one other. It also has the potential to be used over long distances and may be combined with other support for remote collaboration (e.g. videoconferencing).

3.6 Summary

We have created a system that encourages a close working environment by combining portable digital annotation devices with a real-time digital collaboration tool. The iPad hardware and software architecture have been carefully selected to give users their own space to work, while providing them access to the activities of others and also the ability to remotely point out useful pieces of information [12].

4. WHY IPADS?

There is a significant amount of research into collaborative systems and a diverse range of hardware to choose from. Some studies have opted for desktop type systems and others have used the tablet PCs or table top displays. The choice of hardware in a system such as this is an important issue that we considered carefully before designing the interaction that goes with it.

We wanted our system to provide a new interaction design for collaborative working groups. It was central to this vision that the hardware we employed mimicked the size and shape of physical documents. In contrast even to laptops, iPads achieve a size and shape that are small and portable enough to be held like a clipboard or printout. A second factor was the device’s ability to connect wirelessly to the internet, the backbone of the real-time collaboration architecture. Finally, the multi-touch nature of the iPad screen significantly improves the interface interaction promoting almost sub-conscious behavior via ‘lightweight’ gestured techniques [8].

On the topic of “why use iPads?” it may also be useful to consider “why we *didn’t* use eReaders”. While iPads have a similar size and shape to physical books, eReaders also include non-back lit eInk displays - a technology that minimizes eye strain while reading. Despite the obvious problem that most eReaders offer no programmable facility, we have also discovered upon close interaction with these devices, that they also suffer from many serious usability issues.

In summary, iPads are the best hardware choice for our system. The small physical size, and touch screen interaction ensure easy manipulation, and the WiFi facilitates ‘live’ annotation within the group.

5. CONCLUSIONS

This paper has explored a new interaction method for group collaboration. By incorporating the sleek design of iPad hardware with real-time working group annotation, we are able to incorporate the versatility and portability of paper with the instant sharing possibilities of the Internet.

The main features of our new system are:

1. Hardware interaction: iPads are sleek, portable and can be held in the same way as a clipboard and paper;
2. Real-time updating of annotation: notes made on the collaborative document by any user, will automatically be refreshed on every members iPad and will inform them via a text update in the group ‘chat’ box;
3. The ‘look at this’ referencing tool: users can ‘point out’ specific sections of a document to other members of the group which will be placed in the ‘look at this’ queue on the right side of the display.

6. FUTURE WORK

One area of the system that we have not yet considered, is the possibility of interconnecting it with other types of device. Our individual displays may be combined with the shared displays mentioned in Section 1. Morris et al [10] suggest that co-located collaboration tends to include a separate input device for each group member but sometimes incorporates a shared display. This idea could easily extend to our implementation if we could incorporate other devices into the system. For example, as Morris suggests, we could add to the system of iPads, a shared display in the form of projection or table top to add more diversity into the working group and aid in collaboration. Alternatively, not everyone in the group need work on iPads: e.g. one member of the group, perhaps the team leader, may require a larger screen than the other members and will therefore want the system to be executable on a desktop machine instead.

Another area of future study is to test our implementation on several sets of collaborative groups and observe their behavior with this system as well as their usual methods.

The design for this study will start by the selection of an appropriate group of end users: ideally, a group of people who collaborate often either on paper or computers. First then, to form a basis for comparison, we will sit in on one of their regular collaborative sessions and observe their current practices. For example, what equipment do they use? Do they have a single collaborative document or one for every member? How do they make notes? How do they point things out to one another? How many members of the group participate? Are there any dominant members? How many notes were made in total?

Next, we will introduce them to our iPad implementation and sit in on another session where we will observe the behaviors and techniques using the new technology. What we will be looking for here is the differences between the two media. For example, was there a difference in the number of marks made between the systems? Did some users contribute more or less with the new system rather than the old one? Were any problems encountered?

Finally to round off the study, we will conduct a series of individual³ interviews with the participants probing their preferences for each system.

The information gathered from this study will hopefully shed some light on the success of the current system. Firstly, we would hope that users like the idea of the implementation: i.e. does having personal space to work aid in the collaborative process? Is having everyone's marks on the same artifact useful? And was it easy to point out information to others? Furthermore, we would also hope that the system as a whole actually encourages collaboration: i.e. are more notes being made and/or are less dominant members of the group contributing more than they originally did.

This study will hopefully address these issues and perhaps even open up new avenues of future work in the area of real-time collaboration.

7. ACKNOWLEDGMENTS

Jennifer Pearson is sponsored by Microsoft Research and we gratefully acknowledge their support. This research is

³Individual interviews were chosen over focus groups to ensure users with less dominant mannerisms get to express their honest opinion

supported by EPSRC Grant EP/F041217.

8. REFERENCES

- [1] AGOSTI, M., FERRO, N., PANIZZI, E., AND TRINCHESE, R. Annotation as a support to user interaction for content enhancement in digital libraries. In *AVI '06: Proceedings of the working conference on Advanced visual interfaces* (New York, NY, USA, 2006), ACM, pp. 151–154.
- [2] AMERSHI, S., AND MORRIS, M. R. Cosearch: a system for co-located collaborative web search. In *CHI '08: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems* (New York, NY, USA, 2008), ACM, pp. 1647–1656.
- [3] FROMMHOLZ, I., AND FUHR, N. 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* (New York, NY, USA, 2006), ACM, pp. 55–64.
- [4] HINCKLEYSS, K. Distributed and local sensing techniques for face-to-face collaboration. In *ICMI '03: Proceedings of the 5th international conference on Multimodal interfaces* (New York, NY, USA, 2003), ACM, pp. 81–84.
- [5] ISENBURG, P., AND FISHER, D. Collaborative brushing and linking for co-located visual analytics of document collections. *Comput. Graph. Forum* 28, 3 (2009), 1031–1038.
- [6] LINDELL, S. Real-time collaboration tools for digital ink. *J. Comput. Small Coll.* 25, 3 (2010), 24–31.
- [7] MARSHALL, C. C. *Reading and Writing the Electronic Book*. Synthesis Lectures on Information Concepts, Retrieval, and Services. Morgan & Claypool Publishers, 2009.
- [8] MARSHALL, C. C., AND BLY, S. Turning the page on navigation. In *JCDL '05: Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries* (New York, NY, USA, 2005), ACM, pp. 225–234.
- [9] MORRIS, M. R., LOMBARDO, J., AND WIGDOR, D. Wesearch: supporting collaborative search and sensemaking on a tabletop display. In *CSCW '10: Proceedings of the 2010 ACM conference on Computer supported cooperative work* (New York, NY, USA, 2010), ACM, pp. 401–410.
- [10] MORRIS, M. R., TEEVAN, J., AND MARCHIONINI, G. *Collaborative Web Search: Who, What, Where, When, and Why*. Morgan and Claypool Publishers, San Rafael, CA, USA, 2010.
- [11] STEWART, J., BEDERSON, B. B., AND DRUIN, A. Single display groupware: a model for co-present collaboration. In *CHI '99: Proceedings of the SIGCHI conference on Human factors in computing systems* (New York, NY, USA, 1999), ACM, pp. 286–293.
- [12] TANG, J. C. Findings from observational studies of collaborative work. 11–28.
- [13] TUDDENHAM, P., DAVIES, I., AND ROBINSON, P. Websurface: an interface for co-located collaborative information gathering. In *ITS '09: Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces* (New York, NY, USA, 2009), ACM, pp. 181–188.