RWTH AACHEN UNIVERSITY

An Interactive Surface for Literary Criticism

Diploma Thesis at the
Media Computing Group
Prof. Dr. Jan Borchers
Computer Science Department
RWTH Aachen University

by
Stephan Deininghaus

Thesis advisor:
Prof. Dr. Jan Borchers

Second examiner:
Prof. Dr. Karin Herrmann

Registration date: Sep 03rd, 2009
Submission date: Apr 15th, 2010

I hereby declare that I have created this work completely on my own and used no other sources or tools than the ones listed, and that I have marked any citations accordingly.

Hiermit versichere ich, dass ich die vorliegende Arbeit selbständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt sowie Zitate kenntlich gemacht habe.

Aachen, April 2010
Stephan Deininghaus

Contents

List of Figures

List of Tables

Abstract

Professionals in literary studies regularly need to work with large amounts of complex text. Special forms of editions have been developed to make these texts accessible to the literary critic. Nowadays, there exist both traditional printed editions and digital editions for use with standard computers.
In this thesis, we present an examination of the characteristics of these editions, their users, and the users' work processes from an HCI perspective. Based on these findings, we propose a design for a working environment that enables the user to read, navigate, and personalize printed and digital information on one integrative interactive surface, benefiting from the respective advantages of both mediums. This also leads to a promising new way of structuring content in text editions.
Furthermore, we portray the evaluation and refinement of this proposal through two full design iterations using prototypes — each iteration backed by a qualitative user study — and outline possible future work.

Überblick

Bei ihrer Arbeit haben es Literaturwissenschaftler oft mit umfangreichen, vielschichtigen Texten zu tun. Um eine Basis für die Arbeit mit diesen Texten zu schaffen, wurden spezielle Editionsformen entwickelt. Heutzutage gibt es sowohl traditionelle gedruckte Editionen als auch digitale Editionen als Software für Computer.
In dieser Diplomarbeit werden wir zunächst charakteristische Eigenschaften dieser Editionen, ihrer Benutzer und deren Arbeitsprozesse aus dem Blickwinkel der HCI darstellen. Ausgehend von diesen Befunden erarbeiten wir eine Vision einer Arbeitsumgebung, die es ihrem Benutzer gestattet, gedruckte und digitale Materialien auf einer einzelnen, interaktiven Oberfläche zu lesen, zu ordnen und zu personalisieren, wodurch die Vorzüge beider Medien genutzt werden können. Davon abgeleitet beschreiben wir zudem einen vielversprechenden neuen Ansatz, das Material einer Edition mit Hinblick auf die Arbeit in einem solchen Umfeld zu strukturieren.
Ferner dokumentieren wir die Bewertung und Weiterentwicklung des vorgeschlagenen Designs durch zwei volle Designiterationen, die sich jeweils auf eine qualitative Nutzerstudie stützen. Zudem erläutern wir mögliche zukünftige Entwicklungen dieses Systems.

Acknowledgements

First of all, I would like to thank Prof. Dr. Jan Borchers for introducing me to the wondrous world of HCI. Your dedication has always been truly inspiring to me. I would also like to thank Prof. Dr. Karin Herrmann for enlightening me on the field of literary studies and for giving me confidence in my work.

A big "thank you" goes to my advisors, Max Möllers and Moritz Wittenhagen. I highly appreciate the effort you put into this project; always there to provide insight, motivation, optimism, helpful pointers, and constructive criticism, while letting me choose the emphasis of my research by myself.

Thanks to the folks at the Media Computing Group and HumTec. It has been a tremendous experience working with you. Special thanks to Patrick, Simon, and Malte for their support and technical advice.

My apologies to those who had to suffer from my moods during the last months, particularly towards the end of my thesis. I appreciate your bearing with me as well as you could.

Finally and most importantly, I would like to thank my family and friends — especially my parents — for their love and support all along the way. I would not have been able to achieve this without you.

Thank you!

Conventions

Throughout this thesis we use the following conventions.

Text conventions

Definitions of technical terms or short excursus are set off in colored boxes.

PLATYPUS:
The Platypus (*Ornithorhynchus anatinus*) is a semi-aquatic mammal endemic to eastern Australia, including Tasmania.

Platypus

Source code and implementation symbols are written in typewriter-style text.

```
myClass
```

The whole thesis is written in American English.

Chapter 1

Introduction

For ages, technological innovation has transformed our lives. Common knowledge has it that any new technology will first be adopted in professional environments before seeping through to the general public. Information technology is no different in this regard. In the wake of the advent of computers in the late 20th century, many professions saw changes to established work processes to profit from these newly-won capabilities. For some, however, this spelled compromise. Take architects, for example. Although they gained advanced *Computer-Aided Design and Drafting* (CADD) powers from this move, working on their drafts on a regular workstation also became more complex and indirect than it used to be at the drafting table. *Human-Computer Interaction* (HCI) research is still struggling to bring back the ease of interaction that once was.

The humanities have a similar tale to tell. In studying texts, literary researchers rely on editions of these works that are specially crafted to stand up to rigorous scrutiny, provide variants and contextual information, and in general make the facets of the text accessible to examinatiion. Yet by the complexity these very requirements effect, these editions become products — tools, if you will — in which usability is essential to their utility. All the wealth of information in an edition will go to waste if its user is unable to access it. Thus, time and again, editors and typesetters have tested the limits of the printed page to provide as many layers of infor-

mation on a text as possible without sacrificing usability to a point where the edition's utility is degraded. For instance, according to Buzzetti and Rehbein [1998], Theodor Mommsen's 1898 edition of the *Liber Pontificalis* features a wealth of supplementary information that clutters the layout so much that it becomes hardly readable anymore.

There are economic advantages, too; e.g., facsimiles and similar documents are much cheaper to distribute digitally.

Over the past decades, the adoption of computers has led to a whole new kind of edition. These so-called *digital editions* present text and associated documents digitally for perusal on a computer display. The digital medium offers traits that can greatly enhance the functionality of editions. In particular, their flexibility makes it possible to hide additional information out of sight until it is needed; suddenly, the user gets to influence the way the edition presents information, so that he can custom-tailor it to his needs. The computer can also help with previously tedious labor, such as searching for a specific word, or scanning multiple versions of a text for shared passages.

But all is not well. Much like we have seen before, the step into the digital realm incurred a change in interaction quality, as there are fundamental differences between working with physical books and digital text in common document viewing applications. Particularly, confining digital editions to a desktop computer's vertical display and mouse/keyboard interaction renders them remote from desk-centric workflows and removes most of the subtle cues and affordances for navigation and personalization that are present in physical documents. Put briefly, common digital editions help solve one category of access problems, but lead to worse interaction on a more basal level.

It is at this level where we in HCI need to explore new approaches to restore rich and effortless interaction to work processes involving complex documents such as editions.

1.1 Roadmap

In the following chapter, we will present literature on research that is relevant to our problem domain. In Chapter

3—"Initial survey", we will perform an analysis of what makes up the bulk of philologists' workflows and relate the results to similar research in the field. Armed with a clear grasp of what we are dealing with, in Chapter 4—"Design process" we will discuss a design process that works its way from the ground up to formulate a vision of a working environment for literary criticism that integrates the competing traits of current editions. This design will see two iterations as we evaluate its performance through user experiments with increasingly detailed prototypes: one iteration based on a paper prototype in Chapter 5—"Paper prototype" and one using a software prototype in Chapter 6—"Software prototype". We will conclude by summarizing our contributions to the fields of HCI and philology, and point out design alternatives and future work in Chapter 7—"Summary and future work".

1.2 Our background

This thesis was issued by the *Media Computing Group* at RWTH Aachen University [MCG].

Research underlying this thesis was conducted in close cooperation with the *Brain/Concept/Writing* group at the *Human Technology Centre* (HumTec). HumTec is a research environment at RWTH Aachen University funded by the Excellence Initiative of the German federal and state governments, seeking to "foster high level interdisciplinary research between the humanities/social sciences and the engineering/natural sciences" [HumTec].

The benefits of this were twofold: firstly, this allowed us to check in with *Subject-Matter Experts* (SME) whenever needed during the design stages, profit from their insight into the subject, and discuss ideas from a number of different professional vantage points. Secondly, in this way we were able to recruit a number of qualified volunteers from the field — from students to tenured academics — who were not involved with our project and therefore made unbiased interviewees and test subjects.

Chapter 2

Related work

Relevant contributions addressing the outlined problem can be expected to come from both editorial studies and HCI research. Within the field of HCI, we concentrated our survey of the available literature on the areas of document-centric interaction (i.e., navigation, personalization, etc.; for both digital and physical documents) and digital annotation of physical documents, interaction with documents on digital tabletops, and Tangible User Interfaces.

Novel editions

Philological research has resulted in sketches of how information technology could bestow a degree of flexibility on editions that offer novel angles on the text. Such visions exist in the context of editing medieval manuscripts [Carlquist, 2004] and expressing how a text changed over time [Buzzetti and Rehbein, 1998].

Digital text

Whereas philology focuses more on the functional rewards that new models for editions promise, HCI has explored the more basic questions of what interaction with (predominantly) textual information could look like and how it could be implemented. In a sense, this notion dates back as far as the vision of the *Memex* [Bush, 1945].

Today, a wide array of techniques for presentation and navigation exists for digital text [Cockburn et al., 2006, Buchanan and Owen, 2008], including approaches that explicitly mimic physical books [Card et al., 2004, Chu et al., 2004, Liesaputra and Ian, 2008]. Personalization of digital documents has been explored [Hill et al., 1992, Hoeben and Stappers, 2000].

Augmented text

The *Xlibris* system allowed users to annotate digital documents on a tablet with handwriting [Schilit et al., 1998]. Other research covered how pen markings in physical documents and information technology can be combined to enhance functionality while retaining rich interaction. This includes the *a-book* laboratory journals [Mackay et al., 2002], *ButterflyNet* field research notebooks [Yeh et al., 2006], lecture notes [Steimle et al., 2008], and generic documents [Guimbretière, 2003, Liao et al., 2008, Steimle, 2009].

Comparing mediums, interaction fundamentals

Characteristics of reading and annotating digital text and printed text have been thoroughly explored both individually and in comparison [O'Hara and Sellen, 1997, O'Hara et al., 2002, Marshall and Bly, 2005, Everitt et al., 2006]. Recent studies have extended this line of research to novel digital display surfaces [Morris et al., 2007, Terrenghi et al., 2007, Piper and Hollan, 2009]. Fishkin developed a design space for Tangible User Interfaces [Fishkin, 2004].

Tabletops

Other studies look at how horizontal interactive surfaces affect daily work routine [Wigdor et al., 2007, Shen, 2006] and ergonomics [Toney and Thomas, 2006, Müller-Tomfelde et al., 2008]. Han pioneered *Frustrated Total Internal Reflection* (FTIR) touch detection [Han, 2005]. A number of computational models exist that infer gestures and intentions from touches [Hancock et al., 2006, Moscovich and Hughes, 2008, Liu et al., 2006]. Physical objects on a tabletop surface can be identified through means like *SLAP* markers [Weiss et al., 2009]. Furthermore, there has been research aimed at bringing this kind of table-based interaction to a mobile setting [Wilson, 2005, Kane et al., 2009].

Document management on tabletops

Turning to document augmentation in a desk-like setting, the *DigitalDesk* system was one of the most influental systems today that demonstrated digital augmentation of physical documents [Newman and Wellner, 1992]. More recently, the *DocuDesk* system has chosen a similar approach, but focuses on managing links between physical and digital documents on the desk, including the notion of "rehydrating" a previous working state [Everitt et al., 2008]. Koike et al. mixed a printed textbook with illustrative interactive content next to it on the table [Koike et al., 2000], whereas the *Listen reader* demonstrated aural augmentation and page identification through RFID tags [Back and Cohen, 2000, Back et al., 2001]. The *WikiTUI* system allowed the user to annotate physical books with digital notes [Wu et al., 2007]. Rekimoto et al. showed a table-based method to attach digital items to physical objects [Rekimoto and Saitoh, 1999]. The *Starfire* video demonstrated effortless digitization of documents [Tognazzini, 1994]. In early 2010, Hartmann et al. [2010] presented the *Pictionaire* project, which encourages users to mix physical and digital items on a tabletop by enabling them to create digital versions of physical documents on-the-fly.

Tabletop applications in other domains

Other domains have seen augmentation through interactive surfaces as well. These include concept mapping [Oppl and Stary, 2009], urban planning [Underkoffler and Ishii, 1999], and web site information design [Klemmer et al., 2000].

We conclude our search for related literature with the impression that there is currently no research on the support of philological work involving editions from a post-desktop HCI perspective. Handling of regular documents in various settings has been explored fairly well, including work processes that include documents that oscillate between digital and physical representations, work processes involving documents on tabletops, and digital annotation of physical documents. However, we could not find satisfactory research concerning interaction with complex, multi-faceted documents (especially editions). We intend to chart this blank spot in this thesis.

Chapter 3

Initial survey

. . . in which we meet our users, look at how they work, and ask them what they like about it.

3.1 Who are our users?

We identified professionals concerned with the study of literature as our preliminary main group of target users. For the most part, we assumed this to be persons who perform thorough analysis of certain aspects of existing texts — oftentimes referring to one or multiple editions of the text — as well as researchers who specialize in producing these editions.

Furthermore, we conjectured that a new approach to the analysis and presentation of text might also prove beneficial to other fields in which one frequently has to work with complex and large amounts of text. Such fields could include the humanities in general or law.

3.2 Editions

Edition

For their work, literary researchers require an edition of the examined text that meets high scientific standards. In this thesis, for "edition" we adopt the definition of Plachta [1997], stating that an edition is "a reliable text that provides the basis for any historical or interpretative examination" combined with an *apparatus*. Typically, edition projects are extensive scientific endeavors that may run for several years or even decades.

Apparatus

APPARATUS:
In the original sense, a *(critical) apparatus* is that part of an edition that contains variants of the edited text. In a wider sense, by *apparatus* one may mean all those components of an edition that serve the documentation of the edited text, including variants (in the *variant apparatus*), but also descriptive material such as listings of known sources, historical background, *editor notes* that describe the rules by which the edited text was conceived, etc.
In this thesis, we generally subscribe to the latter meaning.

It is an important realization to make that the apparatus is an integral, defining component of any edition, not just a mere addition. Furthermore, the exact makeup of the apparatus depends heavily on (and, in turn, defines) the type and intended focus of the edition. There are even more external factors that can have an influence on the form of the edition, e.g., how much material there is available (and how large a subset of that the editor is planning to publish), how the original author structured his work, in which epoch and under what circumstances the text was created, the condition of preserved documents, etc.

Digital edition

Historically, editions have been published in printed volumes. Nowadays, *digital editions* for use on a computer have been developed, either as an addition to printed volumes or as an all-digital format. These commonly allow the user to customize the way text is displayed and add advanced features such as hyperlinks, full-text search, database integration, and an extended collection of facsimiles, sources, and other material that might be impractical or prohibitively

expensive to distribute in print.

3.2.1 Types of editions

Plachta [1997] also provides a hierarchical typology of editions, the essence of which we will reproduce here.

Historical-critical editions are arguably the most extensive and scientifically rigorous kind of edition. Targeted at professional academic audiences, they feature expansive apparatuses in an attempt to make the *historical* component of the edited text accessible (i.e., documenting the text's creation as a historical process) and to describe the employed process of *textual criticism* that led to the constitution of an edited text (and likewise to the creation of the various related apparatuses). In these editions, there is a clear dedication to providing a scientifically sound basis for the professional to work from, typically incurring a high degree of complexity.

Historical-critical edition

Scholarly editions are a step down in complexity from full historical-critical editions while retaining the same foundation. (In fact, scholarly editions are often directly derived from historical-critical editions.) The goal of these editions is to provide the text for scholarly purposes. Consequently, they do not feature the same wealth of highly specialized apparatuses as historical-critical editions. To aid the analysis of the text, however, their apparatus may take a more pronounced interpretative stance, something that is not appropriate for historical-critical editions.

Scholarly edition

Finally, *reading editions* provide the text with little or no apparatus at all. As the name implies, these are intended for reading alone and are not suited for literary analysis.

Reading edition

In this thesis, we will focus on historical-critical editions, as these are both the most complex type of edition and most commonly used among our target user group. We will, however, introduce mechanisms to customize the apparent degree of complexity of an edition.

3.2.2 Types of variant apparatuses

Turning to the *variant apparatus*, different edition projects have created different ways to provide variants. According to Plachta [1997], we can distinguish between four major methods here:

In computer science lingo, these are basically patch instructions.

Lemmatization: differences between the edited text and the text of a specific variant are indicated in the variant apparatus by listing just the position of each difference (e.g., page and line number) and the difference itself (e.g., a word from the specified line of the edited text and the word by which it is replaced in the variant).

Integrated variant apparatus: not unlike the lemmatization method, variants are given only through their differences from the main edited text; only here, these differences are printed inline with the main text (marked accordingly, of course).

Stair-type presentation: the creation of a text is decomposed into a series of corrective steps, which afford a stair-like presentation in which each such corrective step is set on a separate line.

Synoptic presentation: Variants are reproduced in their entirety side-by-side. This trivially provides full access to the complete text of all variants, but also takes up a large amount of space, which can make it infeasible to present a large number of variants in this fashion on a printed page.

3.3 Initial interviews

The author of this thesis joined the project at a point at which three interviews were already through, and thus only led one interview and participated in two others.

We performed a total of six initial interviews with members of our target audience. All interviewees were literary scientists with professional expertise in the production of scholarly editions, all of them recruited from the university's literature department. Four of them had had little or no prior exposure to our project. We video-captured one interview session as a part of which we had asked the in-

terviewee to demonstrate how he would work on a sample task of textual analysis.

The main goal of these interviews was to get a better understanding of our potential users, their current work routine, and their needs. Accordingly, we designed these sessions as semi-structured interviews in which we would make sure to touch on each of a fixed set of core topics, but would otherwise follow up on interesting points and discuss matters in depth that the interviewee deemed important. We finished most interviews in under one hour each.

Here is a representative list of core topics we discussed during a typical interview session:

- asking the subjects to introduce themselves, tell us with what they are currently concerned
- what does their typical work routine look like, if there is one? If possible, have them verbally walk us through such a task
- what do they like or dislike (and why)...
  - ...about their work in general?
  - ...about working with physical documents?
  - ...about using computers for their work?
- how do they take notes? Do they annotate physical, digital documents? If so, in which manner?
- asking the subjects to imagine a "magical" device to help them with their work — what would that help them with, which properties would it have?

3.4 Refining the target audience

Based on the obtained material, we refined our initial assessment of our target audience to contain two major groups of users:

Note that these two groups overlap, since editors may also refer to preexisting editions in their work.

literary critics, i.e., scientists who routinely peruse editions in order to analyze literature, and

editors, i.e., philologists who are concerned with the creation of such editions.

Due to lack of substantial data, we will not consider the requirements of other potential beneficiaries (general humanities, law, etc.) for the time being.

3.4.1 Personas

For these two groups, we constructed *personas* according to the guidelines set forth by Cooper et al. [2007].

Persona

PERSONA:
A believable description of a fictional person that conveys relevant characteristic traits of a well-defined portion of the target audience.

According to them, the use of personas as a design tool offers a number of advantages:

- the formative processes of distilling characteristic traits from raw behavioral observation aid the designer in *determining users' goals and needed product functionality*;
- personas are a great device for *communicating with designers, developers, and stakeholders*, since their narrative structure makes it easy to relate to requirements and constraints;
- they are a valuable instrument to *measure a design's effectiveness* in an effortless (albeit limited) way.

We constructed a primary persona for the role of a *literature researcher* (cf. Figure 3.1) and a secondary persona for the role of an *editor* (cf. Figure 3.2). We also checked with SMEs

to verify that these personas matched real-world job responsibilities and tasks, and incorporated their advice into our personas.

3.5 Results

Let us now take a look at the findings from the initial interviews in detail. In the following section, we shall group these findings by common themes in the work processes of our users and discuss each of those as we go along. We will discuss how these findings influenced our initial design in the next chapter.

3.5.1 Working with text

The core work processes in the analysis of literature exhibited characteristic similarities across all of our interviewees. Reading and cross-referencing are the predominant activities here. In the analysis of any given aspect of the text, the literary scientist relies heavily on accompanying material from secondary literature, this frequently being historical-critical editions. Furthermore, there appear to be at least two distinctive modes involved in the study of the text, both of which we will need to consider when designing our system.

First off, in certain phases, the scientist will immerse himself in prolonged, mainly *sequential study* of a single (or a few) piece of information. His focus of attention will rest on the document (or the set of documents he selected), and there will be little interaction with other documents in his workspace.

Conversely, one can also expect to witness phases that are dominated by non-sequential bursts of reading activity across different pieces of text, mainly in order to scan for similarities or to follow references in the main text, the apparatus, or personal notes. During these phases, rearrangement of the workspace is common as the user's attention shifts between pages and across documents. Activities in

PRIMARY PERSONA: RESEARCHER

Meet Lisa. Lisa is 35 years old and holds a position as a literature researcher at her university's philological department. Her goal is to understand the creative processes behind the work of an author, and to assess its impact and reception. To analyze a given piece of text, she relies heavily on one or more editions of this author's work, as well as essays and secondary literature by other researchers in the field. Usually, she does not have access to authentic versions of the text; in fact, such versions may sometimes not even exist.

Lisa works with both printed, physical editions, and digital editions on the computer on her desk. Her choice of edition may be influenced by the exact nature of her task. Although a computer user of average proficiency, she feels somewhat discontent with the lack of tangibility and imposing stance that digital editions offer in comparison to traditional printed editions. However, she does appreciate the high customizability and advanced features such as full-text search that come with digital editions.

Let us now follow Lisa along for a bit as she analyzes a poem for its creative context. She starts by placing a stack of books and items on her desk that she thinks will be relevant for this task. Among others, these include a historical-critical edition of the author's works, some further secondary literature, her notes, and pencil and paper for note-taking. She begins by conveniently spreading out these items across her desk, then pulling one or two of them closer for examination. As she reads, she may decide to compare the current piece of text to another passage of text, or to an earlier version of the text. From time to time, she will refer to the editorial commentary to clarify the meaning and background of certain passages, or to gain additional insight.

Over time, Lisa has grown quite accustomed to her books. She has personalized them through markings, underlining, and Post-It notes, so that when browsing through them, she will sometimes be reminded of a different train of thought that occurred to her when she had worked on a piece of text in the past. She considers this aspect of her work enlightening, and would be disappointed if she had to lose it. Indeed, when working on a research paper, Lisa much prefers to jot down short notes on paper as she goes along, rather than put up with a computer. Apart from taking notes, she will also scribble down excerpts from the texts, to be used in her report later on. However, organizing all these notes can become troublesome to her, especially in larger-scale projects, where she may have made a large number of notes.

Quite often, she may even need to tend to other texts than the current one, these shifts of attention being triggered either explicitly (e.g., by a reference in secondary literature) or implicitly (say, when Lisa remembers a similar passage of text in another book). However, in many cases, this means another trip to the library to fetch the referenced text (or assembling a list of books to retrieve later), which she finds rather annoying.

Figure 3.1: The primary persona we constructed for the role of literature researchers.

SECONDARY PERSONA: EDITOR

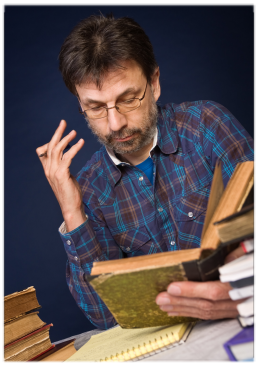

Edward (55) currently works as supervising editor at a renowned research institute that specializes in texts from the late 19th and early 20th century. Over time, he has had the chance to work on several long-running editorial projects, with a wide range of methods and tools. This has reinforced his belief that no two projects in editorial science are alike. Only in recent years has Edward worked on a digital edition, even though he has explored the possibilities of a computer for the organization of his work on printed editions for a long time, and has used them to collaborate with other researchers.

When Edward takes to working on an edition, he first has to assess the collection of textual witnesses, carefully scrutinizing every one of them – what material properties does the witness have, how does the handwriting look, etc. Depending on the type of project, he will need to transcribe individual documents into a more readable format, and possibly provide reproductions of the original. Edward really enjoys this kind of physical access to the material. On a larger scale of things, Edward will also want to infer some kind of relationship among these individual witnesses, such as ordering or grouping them chronologically or thematically. All of this serves an overarching goal of his work, namely to identify influences and to explore, deduct, and document the stages of the creative processes behind the author's pieces. On the other hand, he also strives to capture and represent the reception of the piece by the public (or parts thereof), and examine its impact on other work.

Conceptually, his work mirrors the work of a literary researcher: where the user of the final product will have a specific question about the text, Edward has to make sure that the edition will be able to help him. Where the user will want to interpret the text, Edward has to provide a framework to enable this task, but must take care not to lead into an interpretation with personal opinion. Where Edward knows that other passages of text are related, he has to make this connection visible.

Edward considers himself an expert for the work of the author he edits (and rightfully so). However, at times he may find it unnervingly difficult to foresee the variety of questions his users may have about the text, and how he can strike a suitable balance between providing to little commentary and overwhelming the more casual user with a jungle of comments that are mostly irrelevant for his task, even more so since Edward can provide commentary on multiple levels, from physical properties of the witness over the explanation of isolated words and phrases to in-depth commentary on historical context and the importance of a text.

Figure 3.2: The secondary persona we constructed to represent editors.

these phases can probably be classified best as *navigational* task.

These findings are in line with prior research on the subject [O'Hara et al., 2002]. Unsurprisingly, they also bear a strong resemblance to the famed *spot light metaphor of attention* from cognitive psychology.

3.5.2 Annotation and personalization

This distinction is not entirely rigid; consider for example a *Post-It* note affixed to the text as a kind of non-permanent markup.

Both phases can be interspersed with note-taking activities, where we need to differentiate between markup of the source documents themselves (e.g., underlining or scribbling in the margin) and note-taking in other places (e.g., on a piece of scrap paper, in a text file, or in a private notebook). Among other factors, the mode of annotation depends on the exact document in question, as it is clearly not acceptable to mark up books that belong to a colleague or library. Apart from that, it appears to be largely a question of personal preference.

However, some users mentioned that in the case of notes that are coupled with the annotated document (i.e., margin notes or notes on sheets of paper that are placed between the pages of the document), these notes are susceptible to being "lost" when the annotated document is put pack onto its shelf.

One way or the other, sustained work with physical documents by our users led to noticeably personalized items. Here, we can distinguish between *explicit personalization* (i.e., dog-ears, bookmarks, deliberate markings as described above) and *implicit personalization* (i.e., the way in which a book's pages and spine wear out, grease up, stain, and tear in response to the usage patterns of its owner).

3.5.3 Working environment

By the very nature of the aforementioned tasks, there are some characteristic properties to our target users'

workspaces. As we found out, the user typically sits in a static position in front of a desk, relevant physical documents spread out on its surface. This spatial arrangement of documents is tweaked in varying ways as the subject's attention shifts between documents; see Figure 3.3 for an example of a typical workspace. Such a setup may be complemented by some sort of computer for access to digital documents, access to online resources, and text production.

Figure 3.3: A still frame from one of the interviews, showing a typical configuration of a literary scientist's desk while working on a sample task. One can make out several editions and facsimiles of original texts being used in parallel.

Another characteristic property of the scenarios we evaluated is that the user will typically work alone, allowing him to focus his undivided attention on the task at hand. In other words, collaborative scenarios are few and far between during the actual analysis stage, and we will therefore neglect them for the purposes of this thesis.

cf. Section 7.2.6—"Collaboration" for possible future work on this topic.

3.5.4 Printed vs. digital editions

As stated in the introduction of this section, we asked our interviewees to share their opinion on working with physical

and digital documents, particularly editions.

Common perceived advantages of physical documents were their dependability and long-time persistence (as related to one's personal library), the haptically rich interaction and annotation they afford, their obliviousness to compatibility problems that often arise with digital documents, and their mobility. Users also mentioned the fact that they would grow familiar with their personal library over time as they personalized documents. Additionally, a few interviewees also acknowledged that there was a certain emotionally satisfying quality to the interaction, although the general consensus was that they would use whatever tool suited their current task best.

On the other hand, interviewees frequently complained about the inflexibility inherent in printed documents, the lack of full-text search and referencing facilities, and the relatively small number of facsimiles of historical documents (as the inclusion of too many of these can make an edition prohibitively expensive).

Turning to digital editions, interviewees generally praised their flexibility and advanced capabilities such as full-text search. However, many users stated that digital editions were often cumbersome to use, and that digital documents in general often suffered from the limited screen real estate of current computers, especially when working with multiple documents simultaneously. Finally, some interviewees complained that it was difficult to annotate such documents, or to keep hand-written notes (on physical paper) and digital documents in sync.

3.5.5 Requested features

There were several features that our interviewees commonly requested for the "magical" device that we mentioned towards the end of the interviews. Users frequently wished to enhance their documents with full-text search and indexing capabilities without having to cope with the comparatively poor interaction that digital documents on a desktop computer offer. Others wished for tools that aided them in

handwriting recognition, or allowed them to analyze text at different levels with ease, e.g., by combining data from multiple sources (e.g., databases, other editions) in supplementary layers along with the main text. The wish to superimpose an author's handwriting with a legible transcription also runs along these lines.

Lastly, users would sometimes speculate about sharing and collaboration facilities, both for local and remote collaboration. These features were mainly brought up in the discussion of editorial work processes, i.e., when processing material for publication in an edition.

Chapter 4

Design process

. . . in which we relate our survey results to research in the field, and develop a design to counter our users' problems.

We identified a set of central aspects in our users' work processes for which our design should deliver considerable improvements. Apparently, there is a fine balance to strike between carelessly doing away with tried and proven concepts, and failing to deliver a significant improvement over the status quo. In this chapter, we shall discuss possible solutions to each of these problematic aspects and justify our design choices based on related research and the results of our initial survey. We will follow a top-down approach, working our way from high-level design decisions down to the details. We will conclude by summarizing our design in Section 4.6—"Summary".

. . . or, at the very least, do no worse than the current approaches.

Let us begin by putting forth that from the interview data, we derived two main ways in which a new design could affect our problem domain.

On one hand, we set out to develop *a new conceptual framework for historical-critical editions* (and other complex texts) to present their contents. This is basically about content creation by editors and content consumption by users of the edition.

On the other hand, we have the aspects of *annotation, personalization, and excerption support* in the observed work processes. In essence, this is a matter of both content creation — both explicit and automatically — as well as consumption (i.e., benefiting from personalized items) on the part of the user.

These two groups of factors are undoubtedly related, since specially prepared editions will obviously profit from good support of the latter aspects. However, with a good design, the utility of this second aspect extends beyond novel editions to potentially enhancing work with *all* documents. We shall achieve this by remaining mostly compatible with established processes as to allow the target audience to "ease into" using our design.

This is indeed a fine distinction one should keep in mind, and different points in the design we present below will lean more towards one or the other.

4.1 Workspace layout

Recall that in our interviews, we explicitly explored the perceived respective advantages of working with paper-based documents vs. working with common digital documents (i.e., on a typical vertical computer display). The results showed that when working with multiple documents in parallel or interleaved fashion, amenities like full-text search aside, our target users commonly reported that they enjoyed interacting with a set of physical documents more than they enjoyed interacting with a set of digital documents. We assume that this is at least partly due to two reasons:

Note, however, that higher satisfaction does not automatically imply task appropriateness.

- It is trivially easier to manage multiple documents when one has an entire desk's surface at one's disposal instead of trying to fit everything onto a regular computer screen, where things are bound to overlap quickly. Related research supports our assessment that spatial layout plays an important role in multi-document work processes [O'Hara et al., 2002].

- Rearranging and manipulating items on a computer screen typically necessities going through an extra level of indirection, commonly a low-bandwidth input device such as a mouse or keyboard. We argue that this, too, increases cognitive load and impacts satisfaction ratings.

Low-bandwidth as opposed to direct-touch manipulation.

Another issue arises when both physical and digital documents are needed in one task: since these two types of documents reside on different surfaces, it may be cumbersome to compare documents (e.g., having to hold a book next to the screen, or shifting attention between book and screen repeatedly).

For these reasons, it was clear pretty much from the beginning of the design process that the confines of traditional computer screens and assorted interaction techniques would probably present more of a liability than of an asset for the work processes we target, and it is highly doubtful that current document management and interaction techniques for vertical screens can be reengineered to offset the outlined problems. Instead, we opted to leave vertical computer screens behind and augment the user's desk with a large *horizontal interactive surface* (i.e., a tabletop computer) supporting *multi-touch interaction*. There are several advantages to this approach:

- Due to the increase in available space, the user can apply more natural spatial document layout strategies to manage digital documents much in the way he does with physical documents. The need to shuffle around multiple — potentially overlapping — onscreen windows in multi-document scenarios decreases.
- Through the introduction of direct-touch interaction, we expect to see a higher degree of effortless or even subconscious interaction with digital items on the desk, thus lowering cognitive load to be set aside for these tasks.
- Physical and digital documents now share the same surface, therefore mixing both kinds of documents should put less demand on the user's cognitive resources.

However, ergonomic problems may arise that are either not present or less severe with physical documents or digital documents on a vertical screen. With digital documents on a completely level horizontal surface, reading for a prolonged duration means spending a substantial amount of time hunched forward over the desk. This can lead to discomfort and puts excessive strain on the neck muscles. One way to counter this problem is to tilt the device's surface at a slight angle towards the user [Morris et al., 2007, Müller-Tomfelde et al., 2008].

Optical touch detection

OPTICAL TOUCH DETECTION:
With *Frustrated Total Internal Reflection*, the tabletop surface is flooded with light (typically infra-red), which is kept inside the surface through total internal reflection. When an object touches the surface, the different refractive index at this point causes frustration of the light, which can then be detected by a camera below the table [Han, 2005].
Diffuse Illumination relies on an external light source (mounted below or above the table) that illuminates a diffusion layer on the tabletop's surface. When an object touches this layer, the pattern of diffusion changes, which can then also be detected by a camera.

On the technical side of things, we chose to use a tabletop with back projection and a mixture of *Frustrated Total Internal Reflection* (FTIR) and *Diffuse Illumination* (DI) technology for touch and object recognition. This aspect of the design resembles the approach taken in the *DocuDesk* system [Everitt et al., 2008].

An alternative to back-projection in this regard would have been to go for a top-projection system like Newman and Wellner [1992] or Koike et al. [2000] (also [Kane et al., 2009] in a mobile setting) did; this would have the benefit of being able to project information onto physical documents themselves and making the design less susceptible to occlusion problems caused by physical objects covering digitally projected information on the tabletop surface. However, a top-projection system might be more likely to exhibit occlusion issues due to the user blocking the projection path with his body or other objects, especially when leaning for-

ward over the desk. Nevertheless, we acknowledge that this might be a viable design choice as well.

4.2 Presenting text

Having settled on a desk-sized surface computer as the main environment for interaction with the system, the design challenge that naturally arises from this decision is about how to actually present text to the user. Today, there are basically two different ways to do this, as we have hinted at before: either rely on printed paper, or present text digitally on the tabletop's surface.

4.2.1 Initial approach

In fact, in the beginning we were intrigued by the possibilities a tabletop computer offers for working with fully digital text. We explored the idea of a design that relied primarily on such digital documents, but which, in contrast to documents bound to smaller vertical screens, would allow for free layout in the workspace and synoptic viewing; basically, moving from the mostly window-centric model of common document viewing applications to a system built around a decidedly document-centric metaphor (leaving out the window chrome, toolbars, etc.) — in short, creating a faithful digital counterpart to printed books, but taking advantage of the vast space and dynamic nature of the desk's surface.

This reproduction of printed books even went so far that we decided to keep the page-based presentation of text for our digital model, even though we could have also used the full array of scrolling, thumbnailing, wrapping, and reflowing techniques known to document navigation research. An overview and performance comparison of navigation methods for digital documents is presented in [Cockburn et al., 2006].

Yet O'Hara and Sellen [1997] conclude that the rigid frame imposed by distinct pages functions as an important naviga-

tional cue to our memory when skimming documents, since it enables us to utilize spatial memory to navigate by recognizing landmark features (headings, figures, markings, etc.) in relation to their position on the page. Thus, we argue that sticking with pages as primary containers for text is indeed a sound design decision. This conclusion is supported by the evaluation of systems such as *Realistic books* [Chu et al., 2004, Liesaputra and Ian, 2008].

However, printed documents still hold on to several advantages over their virtual cousins. As we already mentioned in Section 4.1—"Workspace layout", digital documents on a tabletop suffer from suboptimal ergonomic conditions. Furthermore, research concludes that in interacting with paper-based documents, there are *lightweight navigation* processes at work that may be plain out impossible to recreate on a tabletop [Marshall and Bly, 2005]. Printed books also afford mobile use — i.e., away from the tabletop, something this design did not cater for at all — and offer long-term stability; in essence, we left a lot of our target users' gripes with digital documents unsolved. And finally, current tabletop display technology offers nowhere the resolution of a printed page, although that detail may change in the future. All these arguments make a compelling case for traditional printed documents, which is why in the end, we chose to abandon this track.

4.2.2 Second approach

Hybrid documents: printed text for core items, augmented with interactive digital content. Similar ideas have been explored before [Rekimoto and Saitoh, 1999, Koike et al., 2000].

Instead, we decided to pursue an approach that we believe has the potential to give us the best of both worlds: our design relies on *hybrid documents*. These compound documents partly consist of *printed text for core items* that can be used on and off the table. However, when the user places said physical component on the desk, the system should automatically identify it and augment it with *additional interactive information*. For example, an edition might provide one (or a few) designated "main" versions of a text (i.e., the edited text), along with a bit of common information, as volumes of printed text. When the system identifies such a volume lying on the desk's surface, it could augment the main text with additional digital information (e.g., variants, commentary,

personal notes) and afford interaction aimed at navigating and customizing this data, so that the user can tailor the combined experience of the edition to his needs. We will elaborate on this in section 4.3—"Ancillary information".

Whenever the physical object is moved, disappears, or a new object is identified, the clouds of ancillary digital information should follow suit. Therefore, this design basically forces a hierarchy of materialities onto the design, as it promotes physical documents to being the single center of such an information cluster. This inter-material dependency lends our design properties of a *Tangible User Interface* (TUI). One strong point here is due to the fact that research has shown that even when providing similar metaphors, interaction with physical and digital items on a tabletop is quite different. In an experimental study, digital objects on a tabletop appeared to evoke mainly unimanual interaction, whereas physical objects in the same tasks afforded more natural, bimanual interaction [Terrenghi et al., 2007]. Thus, on a cluttered desk where rearranging documents is common, we hope to attain richer interaction with our design then we would expect to see from a design with digital documents as its mainstay.

As a corollary of this newly-won dependency of digital on physical documents, we also get to profit from the rich interaction that physical books afford for large-scale navigation, i.e., quickly thumbing through pages and navigating to an approximate position by judging from the relative thicknesses of stacks of pages. In essence, we transform the physical book into a tangible widget for navigating its own supplementary digital content.

There is also research on how to bring this kind of navigation to digital documents through specialized input devices [Herkenrath et al., 2008].

Note, however, that in this design, we do not in any way banish the existence and use of purely digital documents on the table; in some cases, it will even be a reasonable decision, for instance when a document's physical counterpart does not exist or is currently unavailable. It is just that we strictly *prefer* printed text over digital text.

4.3 Ancillary information

Let us now look at how the supplementary bits of information we mentioned previously are executed in this design. There are two main challenges in this regard: *augmenting the text* and *supporting navigation*.

4.3.1 Augmenting the text

This point deals with the problem of structuring information that supplements the main text. In the context of historical-critical editions or scholarly editions, this could be a full apparatus; several layers of commentary on aspects of interest, references to other texts, or variants. Pushing this kind of information out into the digital material offers a few advantages over printed paper. First, the display of information can be *customizable* — the user may choose to show only commentary on specific aspects of the text, without being distracted by (and allocating space to) information that he is not currently interested in. This contrasts the static nature of the printed page, where this information will always take up space. Second, the properties of the digital surface enable the information to be displayed in a *context-aware* manner — since the system knows about other documents and fragments on the table, be they physical or digital, it can dynamically adjust the display of information to take this context into account. For example, variants of a text could employ local highlighting to make it easier to compare those versions of the text that are currently visible and close to each other, and newly displayed ancillary information could be initially placed on the table in such a way that overlap with other items is minimized.

Another design constraint is that per our survey, we need the system to be able to handle multiple augmented physical documents at the same time. Combined with the fact that we intend such augmenting information to serve a supporting role for a physical "anchor" document, this means that we need some way of making the connection between any digital fragment and its parent document clear. Of course, the obvious choice here is to exploit the gestalt law of spatial

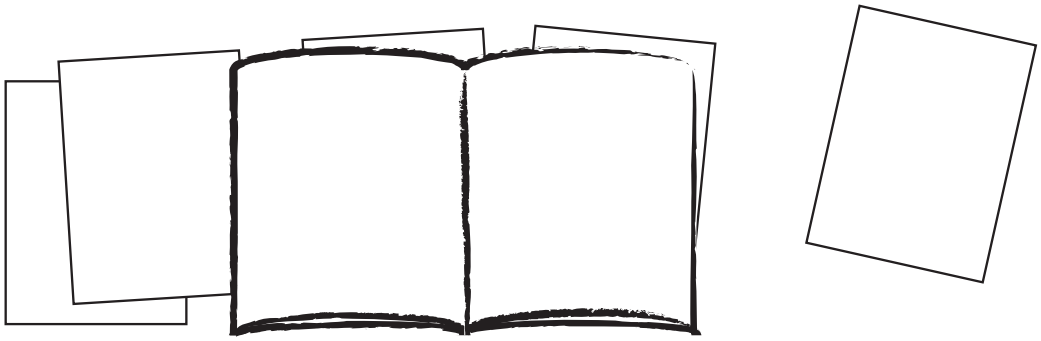

Figure 4.1: A printed book (ragged outline) on top of a pile of digital *fragments* (bits of digital information; represented by clean outlines in this picture) that are related to the currently opened page. The user can grab these fragments by a protruding edge to slide them out from under the book, and push them back under the book when done. Documents and fragments can be rearranged and recombined freely in the workspace.

proximity.

To accommodate these forces, we resorted to a metaphor in which each such fragment pertaining to the printed text (e.g., variants, commentary, etc.) is presented on its own "sheet" of digital paper. We stated that each such *fragment* is initially displayed on the table in a way so that only a small portion of it — by which it is easily identifiable nonetheless — is visible directly adjacent to the physical object to which it is related. The user may then view this sheet by grabbing it by this visible portion and dragging it out from under the parent object. Likewise, expanded sheets can be stowed away by pushing them back under the physical document. The overarching impression we are aiming for with this design is that of the physical document lying on top of a stack of loose sheets of paper, which the user can arrange around the central document in whichever way he sees fit. See Figure 4.1 for an example.

Digital fragment

Where necessary, these individual fragments can also be embedded in a more rigid framework, for example by laying them out in a way that resembles the tabs of a file folder and only allowing to slide them out from under the physical object in a straight line, so that all fragments that share an edge of the object expand in an orderly accordion-like

fashion. We reason that this is a sufficiently useful, yet not too limiting, metaphor. Such a setup is depicted in Figure 4.2. To make the distinction from freely movable digital fragments (as introduced above) clear, we will refer to these bits of information as *slide-out panes*.

Slide-out pane

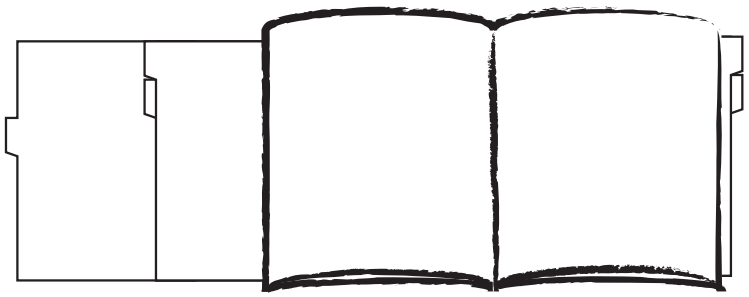

Figure 4.2: A physical book on top of digital sliding panes. The individual panes are constrained to movement perpendicular to the edge of the page, so that they can only be moved in a limited area.

By taking away degrees of freedom and restricting position to the immediate vicinity of the augmented document, we hope to further simplify interaction with these bits of information. In addition, they have a very interesting property in that their contents are always well-aligned with the printed text. Especially in the case of panes that attach to the left and right edges of a document, we can exploit this to provide supplementary content that lines up with the main printed content. For historical-critical editions, this enables us to comment on the edited text in a local and highly customizable way. We hypothesize that this is particularly useful for the presentation of textual variants, as several methods of presenting the variant apparatus (cf. Section 3.2.2—"Types of variant apparatuses") rely on local context. Among these, the method of synoptic presentation stands out. Whereas traditionally, this type of presentation is limited severely by the inflexibility of the printed page and the amount of space it requires for each variant, these problems are much less of a concern in a digital setting. Through the concept of stacks of slide-out panes, the user of an edition can practically resize the effective format of the page to fit a custom selection of synoptic variants. Furthermore, as we mentioned earlier, the system could even compare adjacent variants

to dynamically highlight only these passages of text that differ from each other. (On a side note, literary studies have in recent years seen a debate on *New Philology*, which argues that already the formation of a single edited text is undesirable, proposing instead to provide all versions of the text on equal footing. Synoptic presentation of the variant apparatus obviously suits this stance very well.)

In conclusion, both the approach of collocating printed and digital content on one shared surface and the approach of emphasizing large amounts of available space as a design element really pay off here.

4.3.2 Supporting navigation

One function that digital editions are often touted for is that of *hyperlinks* to additional content. Our proposed design includes tap-activated links adjacent to printed text, or directly embedded in digital content. And while we cannot summon up referenced physical documents on the table, we can do the next best thing: provide information on how to retrieve the document, or possibly display a digital version of the document if available. When such a link to external content is activated, and no digital version of the requested document exists, the system should put a kind of *placeholder* onto the table. We envision this placeholder (or *proxy*) to be a small, free-floating item that fulfills the role of a bibliographic reference that is retained until it is explicitly dismissed by the user, i.e., even when the document from which it was activated disappears. As a bibliographic reference, it needs to carry all necessary information to retrieve the represented document. See Figure 4.3 for a graphical representation.

Harvesting intentions

By virtue of their persistence, they also support a common subtask in active reading scenarios called *harvesting intentions*. Specifically in the line of work our target users do, it is common to read a piece of text with references to other works strewn about, and to collect some of these references for later examination. Traditionally, the reader would have to jot down these references on a piece of scrap paper. With the proxy concept, he can just activate the individual links,

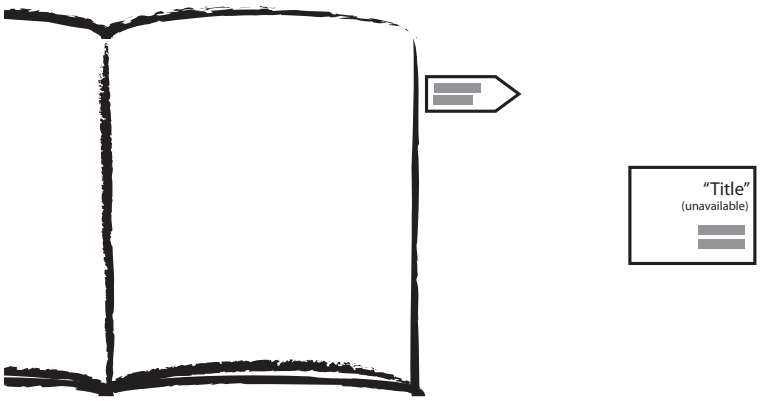

Figure 4.3: An interactive cross reference at the edge of a book (left) and a placeholder for a document that is currently unavailable (right). Such digital cross references are aligned to the position in the printed text at which the link target is referenced. When the user activates the cross reference, a placeholder object is summoned if the referenced document is not currently on the table and no digital version is available.

and when he is finished, there will be a collection of proxies for these referenced documents on the table. The user could then further operate on this batch of items, e.g., print them out, sweep them off the table onto a mobile device, forward them to the local library, etc. A slightly different method to support this process exists in the *Papiercraft* system, which employs annotations of the source document [Liao et al., 2008].

4.4 Annotation and personalization

As we have mentioned before, heavily-used books tend to adapt to their readers' usage patterns over time, and they afford rich annotation and deliberate personalization through markup and other means of modification. In the previous chapter, we already introduced a conceptual division between *implicit* and *explicit personalization* (cf. 3.5.2—"Annotation and personalization").

Digital documents, however, commonly only allow explicit personalization through annotation in a restricted way, typically in the form of keyboard-input text boxes that are placed on or next to the annotated document. This has the drawback that user-provided markings are not as easily identified for navigational purposes. Research does indeed show that under these circumstances, there is a significant difference in annotation behavior between physical and digital documents [O'Hara and Sellen, 1997]. The cited study reasons that this might be because users want their annotations to stand out from the page, in order to avoid the impression that their markings somehow altered the underlying document. The typical contrast between orderly printed main text and handwritten notes serves this division, a fact that is largely disregarded in common annotation facilities for digital documents. (However, turning to implicit or automatic personalization, there has been research to enhance page-based navigation of digital documents with concepts such as *read wear* [Hill et al., 1992] and *electronic dog-ears* [Hoeben and Stappers, 2000].)

Considering these influences, we decided to put an emphasis on handwriting for annotation and explicit personalization in this design. Fortunately, there is technology that can bridge the gap between the physical and virtual realms in this matter: *digital pens*. These behave mostly like ordinary pens (albeit a bit bulkier), but are also capable of capturing and transmitting pen strokes to a computer. The *Anoto Pen* technology [Anoto] we evaluated for this project, for instance, relies on a nearly invisible dot pattern superimposed over the paper, from which it can derive the identity of the page as well as the pen's current position on it. The collected data then constitutes a fully timestamped digital trace of the real-world notes. In the absence of this pattern, the *Anoto Pen* performs just like any other pen.

Digital pens

Armed with this digital trace, we are in a position to retain the user's notes and automatically relate them to recently-used documents. For example, when working with a digital version of a document, software could superimpose markings from the document's real-world counterpart. The *PADD* system used this strategy to keep digital and physical versions of documents in sync [Guimbretière, 2003].

Personal journal

The *PADD* cycle of going from digital to physical documents and back through frequent reprinting is not viable for entire editions, though. Instead, we introduce the concept of an *Anoto*-enabled physical *personal journal* that the user would keep. Similar research on augmented physical journals has previously been conducted in the context of field biology research [Yeh et al., 2006] and laboratory notebooks [Mackay et al., 2002].

In our design, the system would capture all notes the user takes in his notebook. Using the provided metadata from the pen, the system could cluster individual strokes into coherent notes, timestamp these notes, and automatically associate them with documents that are currently known to be in the workspace. Furthermore, when the user places the notebook itself on the tabletop surface, the system should display this information on the table next to the journal in the way it also displays editor-provided ancillary information. This way, when the user pages back through his notebook, the system would be able to indicate the date and context in which a specific note was taken. Recent notes might also float out from under the notebook as individual digital fragments, so that the user can drag them across the table and attach them to other physical documents, thereby creating explicit links in addition to the inferred links. Turning to the other end of this connection, whenever a user places a book on the table that was also present when such a note was taken, that note could be displayed next to the document along with other digital content, even when the containing journal is not currently on the table. Figure 4.4 illustrates this part of the design.

We believe that this design has a number of advantages over other approaches.

- First off, our approach affords familiar, haptically rich input. We consider this to be very important since the issue of interaction quality cropped up multiple times during our initial round of interviews.

- Reliance on a regular pen also affords lightweight input with high expressivity. The user can underline, strike through, doodle, etc. to his heart's content. We

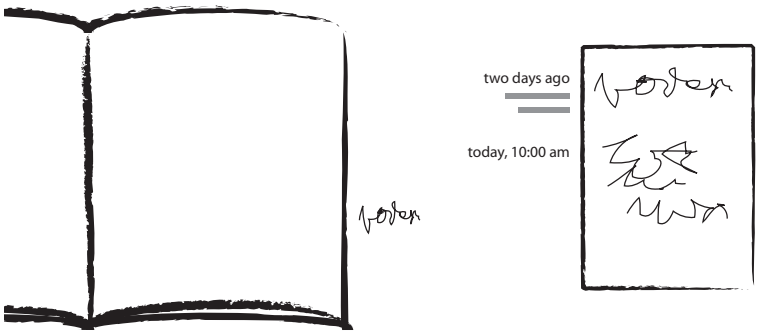

Figure 4.4: A physical book (left) and a physical page from a notebook (right). Digital augmentation is shown near the left edge of the notebook, giving the date and time the note was created together with contextual information. A digital reproduction of the first note is displayed on the table next to the point in the book to which the note refers.

argue that this allows note-taking to be more tightly interleaved with other activities. This is supported by research that compares annotation habits in printed vs. digital documents [O'Hara and Sellen, 1997].

- As we mentioned in the introduction of this section, users show a preference for markings that are easily distinguishable from document content. Relying on handwriting largely achieves this goal. The benefits are twofold: first, this hopefully lowers the reluctance to integrate notes with editor-provided content; second, by making such notes stand out from other content when digitally bound to a document, we expect them to be more likely to act as cues for navigation and trigger recall of related material from memory.

- This design also enables mobile and disconnected use, since notes can be synchronized with the system at a later time.

- By relying on a physical journal, the system gains another level of safety, both technologically and psychologically. This is important because the loss of personal notes are a potentially catastrophic event for a literary scientist.

- This way, personal notes also gain long-time persis-

tence and a certain obliviousness regarding file format issues. Even when the system falls out of favor or the captured data becomes unusable (for whatever reason), the personal notebook will still be there.

- And finally, displaying relevant associated notes whenever a document is placed on the table ensures that notes are still accessible even when the containing journal is unavailable (e.g., discontinued, misplaced, etc.). This mitigates the phenomenon of "losing notes by forgetting them" we introduced in Section 3.5.2—"Annotation and personalization".

4.5 Digitization and excerption

Still, the envisioned blending of physical and digital documents in one workspace can be furthered. We argue that one important step towards seamless integration is to make the barrier between these kinds of documents as permeable as possible. So far, we have touched on the *PADD cycle* [Guimbretière, 2003] as one way to achieve this for short documents or portions of documents. We can, however, introduce a similar mechanism. By positioning a high-resolution digital camera over the desk, the system could be enabled to transparently capture images of any physical document that is placed on the desk. In contrast to the *PADD* model, this solution would be able to capture entire documents without relying on special markings (e.g., an *Anoto* pattern overprint). This approach is largely identical to that used in the *DocuDesk* [Everitt et al., 2008], *Pictionaire* [Hartmann et al., 2010], and, to a lesser extent, *Starfire* [Tognazzini, 1994] systems. Such a facility could complement the design in multiple ways:

There are other options, e.g., using RFID tags [Back and Cohen, 2000].

- Until now, we have silently implied that not only would the system be able to identify physical documents on the surface, but would also know at which page these documents are opened in order to display correct context-dependent augmentation. A high-resolution overhead camera coupled with optical pattern recognition could provide such information.

- In Section 4.3.2—"Supporting navigation" we mentioned that the system could display a digital version of a document until its physical representation (if there is any) has been retrieved and put on the table. A proxy object would then only function as a last resort if neither physical nor digital representation are currently available. An overhead camera as described could greedily (yet transparently) gather such digital versions.
- We speculate that one by-product of transparently captured digital versions is enhanced support for excerption processes, as optical data from the camera can be combined with existing digital content to compile a collection of quite selective fragments from both physical and digital documents for later reference.

4.6 Summary

To summarize, our initial design introduced the following components and concepts:

a multi-touch-capable horizontal interactive surface as the primary working environment to benefit from spatial layout strategies,

physical documents as anchors for a tangible interaction style on said surface,

customizable, context-aware augmentation of these central documents,

hyperlinks and proxies to aid navigation and asynchronous document retrieval,

digital pens for effortless, handwritten annotation of both physical and digital documents, and

an overhead camera for deducing page information and scanning documents.

Chapter 5

Paper prototype

. . . in which we ask users to try out a prototype to identify flaws and strong points of the design.

We evaluated the initial design through a qualitative user study. We will describe the study's design and procedure in this chapter, and discuss the results.

5.1 Study design

For this design stage, we employed a rough paper prototype to evaluate the general applicability of envisioned concepts to our problem domain.

Wizard of Oz interaction

For the experiment, the prototype was brought to life in a *Wizard of Oz* style of interaction, in which the author of this thesis played the part of the Wizard by simulating the system's response to user actions. The behavior of this paper prototype was defined by a behavioral specification derived from our initial design. We include this specification as Section A.1—"Paper prototype behavior" of this thesis for you to review.

We conducted a purely qualitative, between-subjects study with four volunteers, all of them literature students in their

twenties. Three of them had recently taken a class on editorial work, while the fourth student was in the process of preparing an edition for publishing as a part of her thesis. Therefore, all of them were familiar with historical-critical and scholarly editions in general. One student worked for our group and thus had a rough understanding of the project's nature; however, she had never seen the system in action before, or learned much of its behavior. None of the other subjects had previously been exposed to our project, nor had anyone of them previously used a tabletop or other surface computing device, although a few of them were familiar with multi-touch interaction styles such as that afforded by the *iPhone*. Thanks to sparing information given out when recruiting the volunteers, knowledge about the system's behavior was largely under our control and could be specified by our instruction during the test.

We conducted the experiments in a quiet conference room specifically set up for this survey, in order to create a controlled environment. The atmosphere during the tests was generally relaxed, and we took great care to make the subjects feel comfortable about the whole ordeal, stressing that we would not judge their performance, but rather the prototype's capability to adapt to their needs. The entire test was conducted in German, as this was the native language of all participants. Before the core batch of experiments, we validated the prototype and study design through an informal first test run with a volunteer student at our chair. Since the external conditions and the subject's prior knowledge in that session varied from the conditions in our main study (and since that session was never intended to obtain qualitative results in the first place), we will not include the results of that session here.

Validity

Although at first glance, the circumstances and individual tasks in this study (which we will describe shortly) may appear to have little in common with the work processes we observed, we argue that we managed to capture the essence of the targeted processes in this study's tasks, while keeping both the task's and the prototype's complexity down to a level that was acceptable to the subjects (i.e., the subject is not required to spend an excessive amount of time reading, even though this would be the case in the scenarios we target), and afforded agile prototyping and experimentation.

5.2 Prototype components

The prototype consisted of the following components:

- a regular table, designated as the tabletop's surface. Subject and experimenter sat at the long sides of the table, facing each other. At one side, a section was cordoned off with duct tape to accommodate the A/V equipment. Unused parts of the prototype were held in stand-by off the table, out of sight from the subject;

- a fake book with poems, particularly "Fink und Frosch" by Wilhelm Busch, which we altered in a few places to make it harder to understand and enforce reliance on secondary literature. (This poem is a one-page story of a finch and a frog, who, jealous of the bird, engages in a singing contest with him and eventually fails, falling to the ground from a great height.) We chose a Busch poem for their easily readable, entertaining rhymes, which make for a story that can be read in less time than conventional prose. Busch's story-telling is also both witty and, in parts, of almost cartoon-like structure. This decision was also made to help emphasize the trivial nature of the experimental task and support one of our subtasks. Apart from a few more pages of fake, incomprehensible poems and a table of contents, this book was left empty to speed up navigation, and to channel the subjects' attention to the relevant poem;

  . . . or so we thought, anyway.

- a fake dictionary with a number of fantasy words, which was intended to be used by the subjects to understand the Busch poem from the aforementioned item;

- a notebook and pen for subjects' notes and illustrations, in which we provided a sample illustration on the first page to make it look real and give an idea of what was expected;

- two more books (the German constitution and the official soccer rules) unrelated to the actual task, to serve as distractors and add clutter to the workspace.

Those were the "physical" parts of the prototype. To emphasize the division between "physical" and "digital" parts of the prototype, all "digital" items were written or printed on colored paper. In general, this worked satisfactorily (although not perfectly, as we will soon see). Among the "digital" items were: fake commentary (with an interactive bibliographic reference) on the poems, to slide out from under the book; a fake variant of the Busch poem, again to slide out from under the book; various widgets presenting dates and references to physical items, mostly to be placed near such physical items in accordance with the behavioral specification. "Digital" reproductions of the subjects' drawings in the notebook were created as needed during the experiment. See Figure 5.1 to get an impression of the prototype.

Micromobility

MICROMOBILITY:
The small, partly subconscious ways in which we shift a book across the table as we read, or in which we rearrange a sheet of paper when writing.

In order to make the prototype more robust, and to afford micromobility and ad-hoc rearrangement of the workspace, we placed those "physical" items at the center of the interaction (i.e., the poem book and the subject's notebook) on top of transparent slabs, to which we would also attach the stickers that represented "digital" items. This way, when the subject rearranged the books on the table, its digital surroundings would remain in their relative positions. Furthermore, the transparent basis for the poem book sported crude guidance slots for the commentary and variant items, so that the subject could only slide them out from under the book in an allowable direction (i.e., perpendicular to the respective edge of the book) without requiring intervention on the experimenter's part.

The pen we handed to the subject was deliberately chosen to write with green ink, in contrast to the black or dark-blue ink the rest of the prototype used. This served to stress the division between user-created and pre-existing content, since parts of the prototype's interface were hand-written to ease development and emphasize the prototype's unfinished feel. Any reproductions of the subject's notes we created with

Figure 5.1: A part of the paper prototype. Visible in this picture: the main "physical" document for the experiment at the top (white paper) with rigidly attached "digital" commentary panes at the bottom (greenish paper). Note the intentionally rough look and feel to convey that the design is unfinished, and to solicit experimentation and discussion.

a green pen as well, again aiming to clarify the difference between the user's writing and the prototype's interface.

5.3 Experimental procedure

In each experimental session, we followed the same script consisting of an initial questionnaire, a main interaction test with several subtasks, and a semi-structured concluding interview and discussion.

Before the start of each testing session, we explained to the

subjects that we were going to develop a system aimed to integrate with the typical workflow of literary criticism, and that we would like them to work with the prototype of such a system to solve a few trivial tasks.

5.3.1 Initial questionnaire

The form of consent that we made the subjects sign before the start of each session also held a three-item questionnaire, which we used to assess the subjects' prior experience with topics related to our project. In detail, the questions were:

- Have you ever used a digital edition before? (yes/no)
- Have you ever used a multi touch-capable touch screen before (such as that on an *iPhone*, for example)? (yes/no)
- Have you ever used an interactive tabletop before? (yes/no)

One subject reported having used a multi touch-capable touch screen in the past. The same participant also reported having previously used a digital edition, referring to e-books enhanced for the purposes of literary criticism.

None of the subjects had ever used a tabletop computer or comparable user interface.

Apart from the form of consent the subjects signed, we asked them if it was okay for them if we recorded the interaction on the table during the actual experimental stage for later review. We carefully explained to them that the recording's field of vision would be restricted to the actual surface of the table (see Figure 5.2 for an example), and that they could simply choose not to have the interaction recorded. All subjects agreed to have the experiment recorded.

5.3.2 Interaction

We started the actual interaction part of the experiment by telling the subjects that for the purpose of this experiment, the table represented a touch-sensitive surface that was capable of recognizing objects on it and displaying information. We also told them that in the following sequence, we would play the part of the computer, and as such would not be able to answer questions about the system's inner mechanics. We encouraged them, however, to think out aloud if they wished to do so at any time. Furthermore, we once again emphasized that there was no reason to feel pressure, since the focus of the experiment was strictly on the behavior of the prototype, not on how well or quickly they solved a task, and that they were welcome to tinker around with the system.

After that, we handed them the entire collection of physical items (books, notebook, pen) and told them that they could take notes at any time they felt like doing so. We explained that during this experiment, everything on colored paper would be merely displayed on the table's surface, and that they could not pick it up from the table, but would be able to interact with it by touching it, dragging it around, and attaching it to physical objects on the table by sliding it under the object in question. For this rather important part of the explanation, we had prepared a sample digital fragment which we would push around on the table at the same time to demonstrate these concepts.

Then, we gave the subjects a few moments in which we asked them to familiarize themselves with the items we provided. Most subjects examined the pile of books, read their titles, and flipped through one or more of them.

When they appeared to have finished with this, we told them that somewhere in the books we gave them, there was a poem called "Fink und Frosch" written by Wilhelm Busch, then asked them to locate that poem, read it, use the system to solve problems they might encounter, and retell the gist of the story in a few sentences. Basically, the subjects were free to take as much time as they wanted for this task, but all subjects finished within just a few minutes.

This first task was intended as an information retrieval and cross-referencing task. As we noted before, we had altered the poem by replacing a few central nouns with obviously nonsensical gibberish. The poem's optional digital commentary (i.e., a sliding pane attached to the bottom edge of the book) also stated that these words did not really exist, and pointed to secondary literature (which we had also handed over to the subjects) that translated these words to their original meaning. To get at the details of the story, we expected the subjects to discover and follow that connection, as they would most likely come across the prototype's facilities for digital commentary and bibliographic references.

As the last main task, we asked them to imagine that they worked for a publishing house, which wanted to publish a collection of Busch's poems —- "Fink und Frosch" in particular — specifically targeted at children. For this aim, they were tasked with creating two or three simple drawings to illustrate the poem, and had to associate these drawings with the locations in the printed poem to which they referred. A single one of Busch's original drawings for the story was left in the printed book to serve as an example. As soon as the subjects began to draw in their notebook, we quickly reproduced their drawing on a strip of colored paper. When the subject finished drawing for a few seconds, we cut that strip to a size that matched the height of their drawing and attached stickers to it according to the prototype's behavioral specification (e.g., creation time — "a few moments ago" — and automatically guessed associations based on opened books on the table at that moment). Then, we placed these items adjacent to the subjects' original drawings in their notebooks, so that they could slide them away from the notebook and attach them to the poem, in a manner like we had demonstrated at the beginning of the experiment. Figure 5.2 shows a typical scene from this part of the interaction.

We chose to have the subjects create illustrations in this task instead of handwritten notes for a number of reasons: firstly, such illustrations are easy to reproduce by the experimenter (who is looking at the original drawing upside-down), and the legibility of the subject's handwriting does not matter. Secondly, we expected such a task to exert lower cognitive load than a comparable task that would have had the sub-

Figure 5.2: A typical scene from one of the experiments. Subject sits to the right, experimenter to the left of the table. You can see the main book with ancillary augmentation, held in place by guidances underneath the book. The subject in this scene is currently about to create a note, hence the notebook next to the book.

jects formulate written text. Lastly, we expected this to make the students feel more comfortable and emphasize the triviality of the task's direct results once more.

Most subjects asked us if they should continue with the task after drawing and associating two illustrations, at which point we usually cut the task short and declared this part of the experiment finished. We then proceeded to the concluding interviews.

5.3.3 Exit interview

After finishing with the task items, we conducted semi-structured interviews with the subjects to assess the experience the system effected, and invited them to offer their opinion on the study in particular and the general system as a whole. We routinely opened up the interview round

by thanking the subjects for their participation, then giving them an opportunity to share remarks and questions that might have arisen during the test. This segment was usually rather short, as most subjects could not think of anything to say right off the bat. Then, we asked the subjects about their experience using the system, using the following list of questions as a guideline:

- What did you like/dislike about this whole interaction?
- During the interaction, was there a moment where you did not know what was expected of you?
- During the interaction, was there a moment where you did not know how to solve a task?
- During the interaction, was there a moment where you were surprised by what the system did?
- Was there a point at which you would have preferred the system to behave differently?
- What do you think of the supplementary digital material?
- What do you think of the note-taking facilities?
- What do you think is missing from this system to make it useful to you?

As the experiment was designed as a qualitative study, we elaborated and improvised on points of the interaction where the subject appeared to have had trouble or had shown signs of surprise. Our main goal therein was to reconstruct the subject's expectations at that point, and to fathom the mental model of the system that she had developed after such a short time of use. In some testing sessions, these questions were also interspersed with lively discussions about possible improvements (even before reaching the last question of the segment, which was intended to do just that).

Since we felt that annotation support at this stage was still rather rough around the edges, we then took up the chance

to discuss personal note-taking strategies with the subjects, this in a less structured way than before, centered around these questions:

- How do you make annotations? (Do you write in the margin? Highlight or otherwise mark up the text? On Post-Its? On a separate sheet of paper? etc.)
- How do you organize your notes? (Do you use a computer or pen and paper? Is there a hierarchy in your notes? Do you use one continuous document, or one document per topic? etc.)
- How is your note-taking different between books you own and books you borrowed?
- Do you use excerpts when you prepare, say, an essay? If so, at which stage in the process do you use them, and what do you excerpt?

After the discussion, we declared that we were through with our list of questions, and once again offered the subjects an opportunity to ask questions themselves. Finally, we thanked them for their participation, stressed the value of their input to the design process, and offered them to help themselves from a collection of sweets as a little thank-you. We also asked them if they were interested in participating in a follow-up study, once a fully functional prototype was ready.

5.4 Results

Let us begin the discussion of the results with issues related to the study design and prototype itself. First off, several subjects initially hesitated to interact with the prototype. Instead of experimenting, they asked us questions on what would happen if they performed some action. To a certain degree, such behavior is to be expected, and after we stated that we could not answer such questions and encouraged the subjects to just try and see for themselves, the interaction usually proceeded well in this regard. We tend to attribute

this hesitation to the fact that the subjects were unfamiliar with both the nature of the system and the prototyping method, and the fact that we had deliberately provided only the bare minimum of instruction before the test.

Furthermore, in two sessions, we had to explicitly ask the subjects to place the opened book on the table, so that the digital augmentation would be activated. When they opened the book on the page with the Busch poem for the first task (read and retell), they continued to hold the book in their hands to achieve a comfortable reading position. They therefore would solve the task without ever noticing that there was supplementary digital material available. This happened despite us mentioning in the instruction that the table was able to recognize objects that were placed on its surface. This is probably not a serious defect of the study's design, but it suggests that the subjects preferred a reclining reading position to hunching over a book at the height and angle of the table we used for the experiments. Ergonomics definitely seem to be an important factor in this.

In several sessions, subjects initially confused "physical" and "digital" objects on the table. Typically, this became apparent when they tried to place digital objects on top of physical objects, or picked them up from the table for closer examination or rearrangement; the latter especially when moving such items around to the other side of a physical object on the table. In one case, a student attempted to rearrange blocks of text in the printed book when asked to associate her illustrations with the places in the text to which they referred —- trying to make room in the book to fit the illustrations there, just like the given illustration. In the interview, she told us that she thought that the books themselves also had an interactive surface. Thus, it appears that our instructions should have stated that division much more clearly. However, we expected this confusion not to be an issue in the finished system, since the two-dimensional nature of the digital items will be much more obvious there. In general, during the experiment we sometimes were under the impression that the subjects thought the system to be somewhat more akin to a word processor, rather than a device for more reading-centric tasks.

A severe problem with the study itself was that some sub-

jects did not take offense with the words we invented. They thought that these nonsensical words were intended that way by Busch. There are basically two errors we made here: first, to choose replacement words that, although completely fake, had plausible phonetics. Secondly, to choose a poem by Busch, who is notorious for playing with words, in tune with the whole playful feel to his writing. This led to some subjects not referring to secondary literature at all, even retelling the story complete with the fantasy words. Only after we asked them to elaborate on these fantasy words (either directly, or indirectly by asking them to describe the scene in more detail) did they realize that there was more to these words, and started to look for an explanation. Instead of making up words, a much better approach would have been to blacken out portions of the poem, thereby making clear that the text was altered in some way and has to be reconstructed.

A related problem in one session was that the subject had already found the dictionary (secondary literature) that held the true meanings of the fake words when she familiarized herself with the material. Therefore, she was able to finish the corresponding subtask without ever referring to the supplementary material at all.

On the more practical side of things, it was sometimes difficult to keep up with the subjects' actions, especially in faithfully simulating projections on the table when a subject rapidly flipped through the pages of a book. Paper prototyping typically encounters problems when multiple parts of the interface need to change in quick succession. Preparing stickers beforehand and using transparent slabs as the basis for books and their digital augmentation helped somewhat in this regard, but could not completely alleviate the problem. In order to not disrupt the subject too much in her work, and to avoid having to repeatedly call time-out on the interaction to change a minor detail, we decided to err on the side of seamless interaction, which means that we would risk missing a few details here and there during sequences of rapid interaction, but would still catch up as soon as the system returned to a stable state (i.e., the subject reading or weighing her options).

Finally, even though the study tasked the students to make

themselves familiar with the material, think aloud, and feel free to experiment, system coverage was generally surprisingly spare. For example, none of the subjects slid out the horizontal (variants) tab in the first half of the test. Interaction always appeared to be heavily directed at solving the task at hand.

All session finished within the projected 40–45 min timeframe.

In the following sections, we will now discuss the observations we made concerning individual parts and concepts of the design.

5.4.1 Workspace layout, occlusion issues

We observed surprisingly little rearrangement of the workspace by the subjects. This includes large-scale relocation of individual objects and piles (such as organizing books and pushing objects around as the attention shifts from one book to another) as well as micromobility. The distribution of the books we gave out seemed to be rather static, and depended largely on the way the books were strewn out or piled up on the table immediately after we handed them over.

However, all of the subjects did frequently run into occlusion and collision issues, particularly when they summoned horizontal slide-out panes (the variants apparatus colliding with neighboring books), when taking notes in their notebook (digital reproductions and time/association markers colliding with a neighboring book), and when expanding the slide-out pane at the bottom (the commentary almost running off the near edge of the table). In these cases, all subjects automatically set out to rearrange the occluding physical items when they realized that the newly appearing items would appear beneath them. It remains unclear whether they did this to be able to read the newly appearing items, or if it was solely an act of courtesy towards the experimenter when he had to lift up the books to slide the new stickers beneath them. Furthermore, one should consider that this having to briefly lift items off the table certainly

drew undue attention to these items. In a real system, the absence of this effect might likely diminish, or vanish altogether, should a digital fragment be fully occluded by a physical object on the table.

5.4.2 Slide-out panes

In general, sliding panes appeared to work pretty well. All subjects (correctly) assumed that these tabs sticking out from under the book would let them access additional information. Most tried tapping the visible part of the tab to activate it (making it slide out to full extension, according to the behavioral specification), while one subject grabbed the tab and dragged it out manually. Although possible per the specification, no student attempted to extend one of the panes further from the book than necessary to view its contents in their entirety.

Students routinely exposed only the bottom (commentary) pane, as they considered it unnecessary to access variants for their task. One student uttered minor surprise upon finding a variant apparatus here, saying that she would not expect to find such in a reading edition.

In the discussion, one subject remarked that she did not like the sliding pane to extend from the bottom of the book, as it might collide with the near edge of the table, and that she preferred such panes to extend horizontally or to the top of the book. Indeed, such collisions with the near edge of the table did occur, although students then typically just shifted the book around on the table to make room for the pane (see also the section on workspace layout and occlusion issues).

Furthermore, the specification used for the test mandates that when the bottom pane is slid out for one page, the pane for the facing page becomes visible as well. Similarly, when one pane is slid back in, the other automatically follows suit. This was clearly something that the subjects did not expect, even more so since the panes were visibly distinct and had separate handles. Even worse, this coupling frustrated one subject, who attempted to dismiss the commentary for the unrelated poem on the facing page, but thereby also hid the

information in which she took interest. This detail obviously needs to change.

The link widget in the commentary — a bibliographic reference to secondary literature — was well understood. The subjects had no problem understanding what the widget was for (i.e., somehow give access to secondary literature) and how to activate it (i.e., by tapping it). All subjects who found the commentary also activated said link.

5.4.3 Proxies

The concept of using proxies (small, free-floating items on the table's surface that hold bibliographic information about literature that is not currently present on the table, and serve as placeholders and reminders until either that literature is retrieved or the proxy is dismissed) was generally well understood, yet appeared to lack functionality. Whenever a proxy appeared for the first time as the result of the subject following a link in the digital portions of the prototype, they expected it to convey further information when interacted with. Typically, the subject read the text it displayed, than tapped it, appearing puzzled when nothing happened in response. In the review sessions, subjects expressed their expectation at that point that this touch would yield one or more of the following (answers from all subjects compiled into one list, the first two items of which were the most common):

- access to a full digital version of the linked document;
- an interface to order the linked document from a nearby library additional information and meta-information related to the document, such as its ISBN, its table of contents, abstract, or an abridged preview of the document. According to the subjects, this would serve two needs] firstly, to gain additional information that helps with retrieving a copy of the document, and secondly, to perform information triage, i.e., to assess if the referenced document is likely to contain information that is relevant to the reader's current task;

- possibly scans of the document, if it had been on the table in the past.

In one case, the subject remarked that she would prefer the proxy to open as close as possible to the point of interaction from where it was summoned. (We had obeyed the specified behavior to the letter, that is, to search for an open space on the table in a straight line outwards from the initial point towards the far edge of the table. In this case, this put the proxy at quite a distance from the originating link.) In general, though, the connection between the activation of a link by the subject, and the appearance of the proxy appeared to be sufficiently clear, regardless of its actual position.

5.4.4 Annotation support

The part of the prototype which captures and tags handwritten notes (and allows the user to associate such captured items with relevant passages of text) did not work very well in the experiments.

First off, we ran into some issues with the experimental design. Quite noticeably, all subjects hastily flipped past the sample entry we left on the first page of their notebook with no more than a short glance at most, even though we had left ample space beneath it for their own illustrations. These entries were fully set up with fake timestamps and associations, intended to give the subjects a first impression of how items in the notebook would eventually turn out. Since they never saw action during the tests, they were rendered useless. We suppose that this happened due to the subjects' quickly categorizing this drawing as not belonging to them, making them avoid it. Furthermore, resulting sketches were generally much more detailed than we would have needed them to be.

All in all, some subjects showed slight surprise when they saw a copy of their drawing appear on the table, but all of them recognized it as "their" sketch. When we later asked them about this moment, all of them correctly speculated that the system had somehow automatically captured their

writing. (One student was already familiar with the *Anoto* system and pointed into that direction.)

The moment in which the reproduction of the first sketch appeared was a prime source of occlusion issues. To clarify, consider that the poem we used sat in the left half of a pair of facing pages when the book was opened, and that the subjects, all of them writing with their right hand, placed their notebook the the right of the book that contained the poem, so that they could read and draw without rearranging one of the two items. Consequently, the digital versions of their sketches, which were scripted to appear at the left side of their notebooks to avoid collision with the hand that held the pen, surfaced partially under the book instead. After the first such incidence, all subjects created a larger gap between the book and their notes.

Now for the action of associating individual notes with portions of the text. Interaction details in this part of the experiment were sometimes rather unforeseen. For example, one student chose not to slide the digital items under the side of the book, but only let them hover near the edge of the page. Another student tried to attach her drawings to the top and bottom of the page, instead of the left and right edges. When asked about this, she reasoned that her illustrations showed the start and the ending of the story, and that she therefore wanted to emphasize the chronological order of the individual parts. As we already described above, one subject attempted to rearrange the blocks of text in the printed poem, in order to fit her drawings right onto the page. Furthermore, several students attempted to pick up a drawing from the table to move it from their notebook to the poem. Without leaving the table's surface, they would have had to drag the digital drawings in a curved path around to the other side of the book.

At the beginning of the experiment, right after we gave our instructions, several students asked to clarify if they would be able to write directly on the table's surface, or in the book. We told them that they would be able to write in the notebook and in all other books, but not directly on the table's surface. In the exit interviews, it became clear that many subjects would have preferred a different annotation method. Several subjects noted that they wished to just

write on the table next to the poem (in a kind of digital column dedicated to that end), and that the current system of dragging written annotations from their notebook the the poem felt clumsy to them. One subject remarked that with the current design, she would not be able to comment on the digital portions of an edition.

Several subjects also expressed dissatisfaction with the granularity at which notes can be linked to the text in the tested prototype. They argued that they wanted to be able to highlight or comment on individual phrases or words, something the current design does not readily permit.

In the interviews, we also discussed their note-taking strategies with the participants. It turned out that all of them rely on some sort of external document when working with text. This document (or multiple documents) are usually text files on a computer, with one document per task or poem they work on. Subjects were divided on marking up the source document itself. Of those who did, markings were usually restricted to underlining and short comments. One subject said that she avoided such markings out of respect for the book, but would like to write on the page itself if such markings were fully reversible. None of the subjects relied heavily on Post-It notes, as they were easy to misplace. One student did use them, however, as a kind of bookmark (i.e., not left in the book permanently, but only to point to a few selected places in the book).

5.4.5 Wrapping up

We consider the experiments a success, as they shed light on the design's flaws at this point. Furthermore, all subjects volunteered to participate in a follow-up study with a more advanced prototype, expressing great interest in the future of the project.

Chapter 6

Software prototype

. . . in which we translate the design into a software prototype to lay the foundation for the envisioned system.

6.1 Scope and design

We proceeded to develop a second prototype (nicknamed APPARATCHIK, referring to an edition's apparatus) as a native *OS X* application written in Objective-C to run on actual digital tabletop hardware. Due to time and technical constraints, this prototype contained a subset of the functionality we outlined in Chapter 4—"Design process". The behavior of some aspects of the design was slightly changed to reflect findings from the paper prototype evaluation.

We will begin the description of the prototype with a description of the aspects we covered in this iteration.

6.1.1 Ancillary information

The prototype featured ancillary digital information for specific pages of a physical book as described in Section 4.3—"Ancillary information". Both detachable, freely movable

pieces of information (termed *fragments*) and *panes* (bits of information that are rigidly coupled to a specific page and can only be moved along one axis) were implemented. Sliding panes can be expanded by dragging them along their path, or toggled between fully expanded and fully collapsed states by tapping any portion of them. If the user drags a pane and releases it close to its points of full or minimum expansion (in the user study, "close" meaning within a distance equal to 33% of its full range of movement), the system plays an animation to snap it into the respective extremal state.

Concerning detachable items (*fragments*), the system keeps a list of "required" and "user-attached" items for all pages of any document. Therefore, the user can take a fragment from one document and attach it to another page in the document, or even to a page in another document. That item will then remain attached to that document until it is removed again by the user. The user can also attach a fragment to a physical document by picking up that document and placing it on the table so that it intersects with fragments that are not currently attached to any other physical document. Meanwhile, if the source document "requires" the presence of an instance of that fragment, the system will automatically spawn a replacement upon noticing that the original item is gone. (Note that we changed the details of this aspect halfway through the user study, as explained in that section.) On the other hand, the system will ensure at all times that no two duplicates of the same fragment can be attached to a single page simultaneously. Although not directly specified in the original design, we decided to impose this limitation to avoid clutter. Whenever a fragment is newly attached to a physical document, the system will animate that fragment's bounds by expanding and contracting a few times in rapid succession as to provide feedback that the state changed.

The system also needs to ensure that no fragment is ever completely obscured by a physical document, as this would render it permanently inaccessible. For example, this situation might arise if the user places a physical document on the table so that it completely occludes that item. If such a situation is detected, the system will relocate the occluded item towards the top edge of the occluding document, animating the process to draw due attention to it since the user

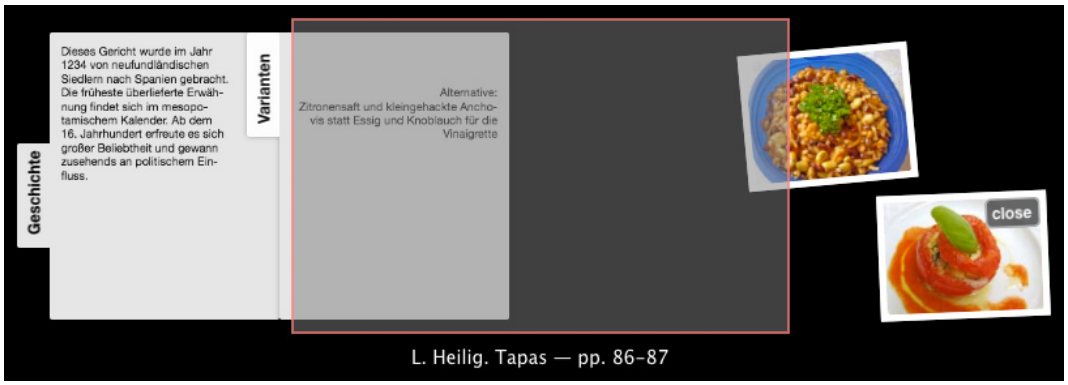

Figure 6.1: Screenshot demonstrating ancillary information. During actual use, the corresponding book would lie exactly on the red frame in the center of the picture, thereby occluding anything beneath it. You can see two panes expanding outwards from the left edge of the book and two detachable image fragments near the right edge. The lower one of these images has been moved so far away from the book's position (i.e., the red frame) that it is now independent from its containing document, hence the "close" button. The line below the book's position signals that the document has been recognized correctly.

may want to correct this attachment.

Apart from these points, this class of items follows a manipulation model that is frequently seen in tabletop applications: the user can drag the item across the table using a single, continuous, moving touch. When two or more concurrent touches are connected to one such item, and a certain threshold of relative movement among the touches is exceeded, this will result in a change of scale and/or a rotation in addition to the translation component. For scaling, we imposed arbitrary lower and upper bounds to limit mishaps due to two moving touches in close proximity to each other. Under no circumstances is the user able to skew the item or change its aspect ratio.

See Figure 6.1 for an example of these items in action.

6.1.2 Navigation

The prototype also covered cross references for inter- and intratextual navigation, along with document proxies for

cases in which no digital version of a document is available and the user needs to be prompted to retrieve a printed representation of a document. The system is intelligent enough to coalesce multiple references to the same document into one placeholder object and to prevent duplicate references to a unique position within a text. Whenever a reference is added to an existing placeholder, the system plays an animation to draw the user's attention to the respective UI element and signal that the user's action did indeed have an effect. Furthermore, when a previously absent document is retrieved and placed on the table, the system automatically converts the previously independent proxy object into a signpost-like message attached to the top edge of the document. The initial design and the behavioral specification of the paper prototype did not specify this. See Figure 6.2 for an example.

6.1.3 Excerption

Full-page excerption functionality was implemented, so that upon request, the prototype appears to produce a scan of an entire page of any physical document on the table. The user can trigger this action by tapping an on-screen button positioned near the top outward edge of any physical page (i.e., for a regular book with facing pages, there are usually two buttons, one for the left-hand and one for the right-hand page; for a single-page document, there is only one button). Excerpts start out as partly tucked under the physical page from which they were created. Note that this deviates from the specification used in the paper prototype, where the user was able to excerpt selected portions of a printed page. This difference is mainly due to greatly reduced implementation complexity of full-page excerption, though. These excerpts then behave just like any other digital fragment on the table, i.e., the user can move them around or attach them to any physical document. See Figure 6.3 for an example.

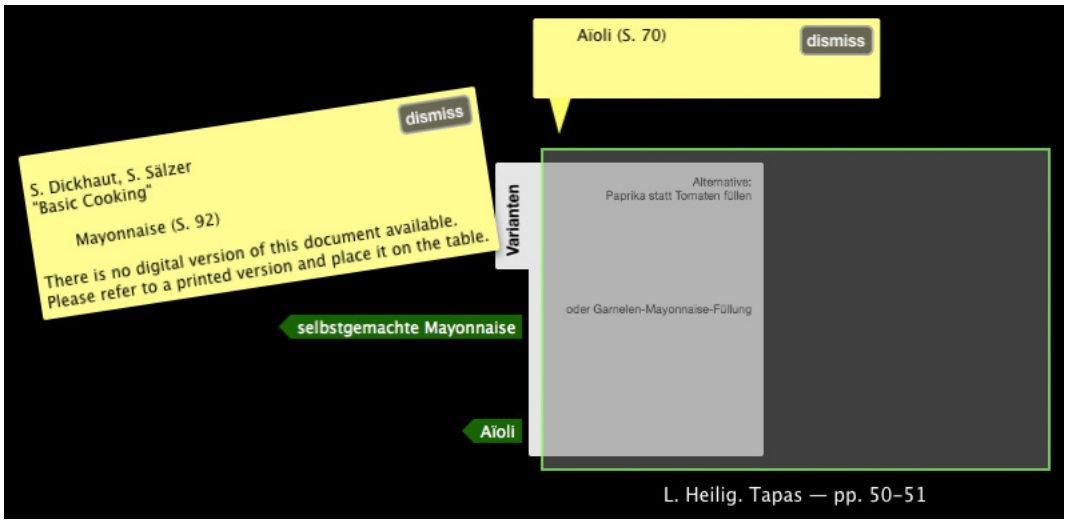

Figure 6.2: Screenshot of navigation facilities. Depicted are two cross references to another position in this document and to a position in another document, respectively (i.e., the green arrows at the left edge of the document). As both cross references have been activated prior to taking this screenshot, two additional UI elements have been spawned. The yellowish box above the document reminds the user of a position in the current document that he marked for review, while the disconnected yellowish placeholder object off the left edge of the book prompts the user to retrieve the referenced document, place it on the table (at which point the placeholder would attach itself to the top edge of that document just like the box above the source document in this screenshot), and open it at a specified page. You can see that both elements have a "dismiss" button in the top right corner in case the user decides not to go through with the navigation.

6.1.4 Annotation

The prototype already provides an extensive infrastructure for pen input, tagging and clustering of captured strokes into higher-level structures, and the ability to display such information along with the document. However, we did not come around to a thorough evaluation of this aspect.

6.2 Implementation details

The software prototype utilizes the *MultiTouch* framework developed by our group for optical touch detection and

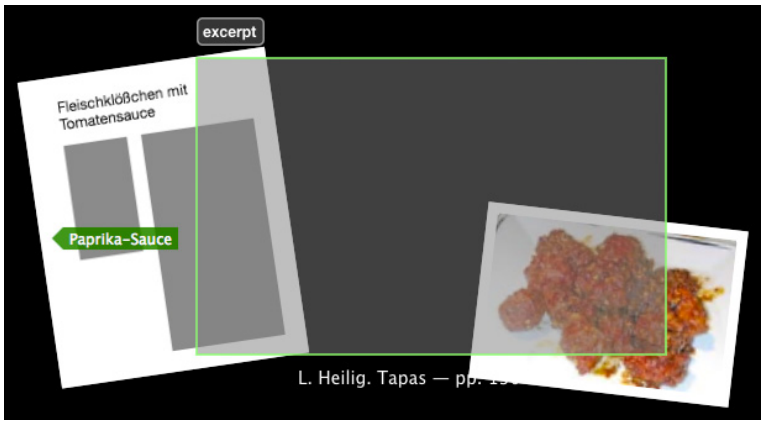

Figure 6.3: Screenshot of the software prototype demonstrating excerption functionality. Note the "excerpt" button above the top left corner of the document, activation of which summoned the digital page excerpt that is shown tucked under the book's left edge in this screenshot.

notification. In the implementation for this thesis, touch events were delivered by the *MultiScreen* event server of the framework, but in principle, one would also be able to use the original *MultiTouch* agent for input (in fact, the system was initially developed against the *MultiTouch* agent). *Core Animation* is used to render most of the user interface.

Architecturally, the software follows a fairly standard pattern for Cocoa applications. Central components are:

- `ApparatchikAppDelegate`: the *application delegate* takes care of proper startup and shutdown, and registers the application with the *MultiScreen* event server. Additionally, the application delegate handles those parts of the user interface other than the main view (i.e., menu commands and configuration windows).
- `MainView`: the *main view* provides the basis for the interface. Specifically, it initializes and manages the *Core Animation* layer that acts as the root of the main interface. This view is also responsible for routing all incoming touch events to their respective handlers. Further duties include reacting to drag-and-drop events, reacting to keyboard events, and switching the view

in and out of fullscreen mode.

- Interactive parts in the user interface (e.g., books, fragments, panes, cross references, etc.) are generally engineered to employ an MVC-style separation of concerns, with the "view" and the "controller" parts combined into one object. For example, *panes* (bits of digital information that augment a physical document) are represented by a tandem of a `PaneModel` class providing general information (i.e., the name of the digital document from which the pane's content can be retrieved) and a `PaneController` class that handles the actual user interface for a `PaneModel` instance by rendering appropriate content and responding to events.

- `BookModelDataSource`: this class constructs a set of book model objects for use in the prototype.

- Further groups of cooperating classes exist for the domains of touch event processing (`MultiScreenNetworkServer`, `TouchTracker`, and others), pen-based annotation, and general utility functions.

6.2.1 Touch processing

We will now describe in more detail how touch events are handled throughout the system, as we deem this both non-obvious and essential to its behavior.

Whenever the *MultiScreen* agent identifies a spot on the tabletop's surface from the camera input, it builds a corresponding `MTTouch` instance and notifies its clients of the change through the distributed notification center mechanism of OS X (`NSDistributedNotificationCenter`).

Client-side, these events are picked up by the application delegate's instance of the `MultiScreenNetworkServer` class. By keeping track of identifying fields of all received touch notifications, this object can then synthesize and update exactly one `MTTouch` object per active touch to present to its clients. The main advantage of this uniquing step

is that it greatly simplifies keeping track of touches in the following layers.

At startup, the application delegate will have set up its `MultiScreenNetworkServer` instance to relay touch events to an instance of the `TouchFilter` class. In this step, additional bits of information are computed and stored in the touch events — mainly timestamp and tap count data that is not currently provided by the *MultiScreen* agent we used during development. As it is likely that this will be corrected in future versions of the agent, this filter was designed to be easily extracted from the software once the necessary data is present on incoming touch events.

Touch events are then passed on to the application's `MainView` instance, which needs to discern three distinct types of events:

- When a new touch begins, the view performs hit-testing on its root *Core Animation* layer to retrieve the `CALayer` instance at this point. The layer hierarchy is then traversed upwards towards the root layer in search of a layer with a KVC key named "touchTracker". If such an entry is found, its value is considered to be a reference to an object conforming to the `MTTouching` protocol (typically a `TouchTracker` object), and the current touch event is relayed to this handler. Furthermore, the main view will associate this touch with the found handler, so that the handler can be notified of further events involving this touch.
- When a touch ends, the view notifies the associated handler (if one was determined at all when the touch began) and ceases tracking this touch.
- When a touch is updated, the view notifies *all* active touch handlers of this change, even if they are not directly responsible for this touch. By notifying handlers in such an order that handlers for layers further up in the layer hierarchy are updated first, we generally attain the goal of "physically correct" coordination of multiple simultaneous touches within a layer hierarchy, even when these touches are handled by independent touch trackers. Such an ordering can be safely

Note that since the containment graph is both directed and acyclic, we can guarantee that a topological ordering exists.

derived from any topological ordering of the graph of sublayer/superlayer relations.

When an instance of the `TouchTracker` class is then passed such a set of beginning/updated/ending touches from the main view, it determines the current location of each touch and converts it to the coordinate system of a specified `CALayer`'s superlayer (usually a layer that contains all other parts of a controller's user interface; assigned during construction). By comparing the last known location (in local coordinates) of every touch to its updated location (again, in local coordinates), the tracker can then compute the matrix of an incremental affine transformation that, when applied to points in the reference coordinate system, would keep these points aligned with the touches (or at least provide a good approximation in the case of three or more simultaneous touches, since we want to avoid skew components in the matrix). The tracker object then passes this transformation matrix to its delegate.

Typically, the tracker's delegate is set to the controller for the manipulated object itself, which also manages the `CALayer`s that make up the interface. Therefore, the controller can simply append the passed incremental transformation matrix to the main layer's `transform` property to update the interface. In cases where the change to the rotation, scale, or translation of the manipulated object must obey constraints (e.g., restricting panes to very limited movement at the edges of a physical document, or imposing lower and upper bounds for the scaling of documents), the controller can directly modify the transformation matrix (e.g., to eliminate scaling), or extract only certain components (e.g., translation along the x-axis) from it.

6.2.2 Hardware

Future iterations of the software are intended to be run on a digital tabletop that members of our group built in parallel to the software prototype. We are confident that this table will provide a reasonable working environment for our target users. This tabletop features a fully interactive area (i.e.,

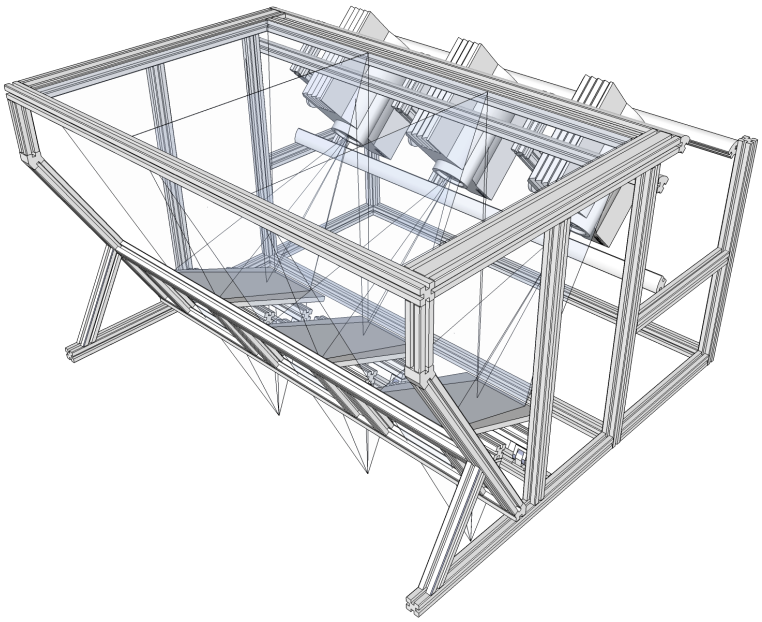

Figure 6.4: Schematics of the final hardware setup. Note the three projectors at the far end of the table, each carefully aimed at a mirror at the bottom of the frame so that its reflection covers exactly one-third of the entire surface. Cameras are mounted directly between the mirrors (not shown in this drawing).

both responsive to touch and capable of displaying the user interface) of approximately $135\,\text{cm} \times 80\,\text{cm}$. Three projectors are linked together to display the interface on the table surface via slanted mirrors below the table. Using this setup, the system can provide a native resolution of 3240×1920 pixels across its surface. Two upward-facing cameras are mounted between the mirrors at the bottom of the device to capture refracted infra-red light from touches. See Figure 6.4 for details. Please take note that even though the construction of the tabletop hardware as described here was relevant to the development and evaluation of the software prototype, the author of this thesis was largely uninvolved in the design of the hardware. Therefore, it does not constitute an original contribution on the author's part and is described here merely for the sake of completeness.

6.3 Evaluation

As in the previous iteration, we conducted a user study to measure the prototype's performance. We decided to do another qualitative evaluation that paralleled the design of the previous user study.

We recruited a total of six volunteer students for this experiment. Two of them were from literary studies, the other four were students of computer science. One subject had already participated in the paper prototype study, while the others had had no prior exposure to the project.

Since it became apparent during the preparation of the experiment that the intended final hardware as outlined in Section 6.2.2—"Hardware" might not be sufficiently operational in time for this round of tests, we decided to use a different digital tabletop at our group for this study. As a consequence of this decision, we had to evaluate using a much lower resolution (a single projector at 1024×768 pixels) and smaller working area, although there still was ample space for two regular-sized books on the tabletop at the same time.

Furthermore, since automatic detection of physical books' positions and opened pages were not yet implemented, we decided to fake this part of the interaction through another *Wizard of Oz*-style of study. To this end, we instrumented the software prototype to respond to keyboard input to place the digital counterparts for certain documents on the table, respond to turned pages, etc. A seemingly good selection of keyboard shortcuts and a relatively fixed script for the individual subtasks of the experiment enabled us to respond quickly to most user actions.

6.3.1 Study design

The low resolution of the used surface in combination with the heterogenous makeup of our volunteer group prompted us to target this study at how well a number of recurring tasks from literary criticism are supported by our design. To this end, we decided to prepare a number of simple tasks

centering around cookbooks instead of regular works of literature. This approach has a number of advantages:

Validity

- First and foremost, we can construct tasks that mirror the work processes we are interested in well, and thus bear significance for our target users. We will argue for a number of such analogies shortly.
- It is relatively easy to come up with a believable setting, believable additional material, and believable tasks if the study requires it. Due to a cookbook's instructional nature, triggering certain actions (e.g., referencing other recipes) does not feel arbitrary.

Sequential reading is a major component of our target users' work routine, but we are not trying to address raw reading performance.

- The books we used had a very sparse layout with lots of whitespace and pictures, and text was generally of low complexity. This meant that the subjects would spend less time scanning/reading the page than with many other types of documents.
- Individual recipes impose clear boundaries in the provided material. Thus, the subject is unlikely to stray off task and go on examining material that is not relevant to the experiment (something that happened a few times with the paper prototype).
- Since we had to plan for wildly different proficiencies regarding literary tasks among our participants, we hoped to create a common ground for them with which they would all be equally familiar.
- For some bizarre reason, many cookbooks are less expensive than blank paper. This allowed us to experiment freely with a number of different books.

At the beginning of each study session, we gave a short introduction and asked the subjects to fill out the form of consent and a questionnaire to measure their familiarity with various interaction technologies. Supposing there were no initial questions on the part of the participant, we then started with the first of three task.

In the first task, we handed the subjects a book with tapas recipes and told them to try and use the system to retrieve

as much information as possible on a specific dish contained therein. As soon as the book came to rest on the table with the indicated page open, we made the system display additional data for the current page around it and observed how the subject interacted with this. We provided two sliding panes that contained a variation of the recipe and some historical background, respectively. Two detachable images of the dish stuck out from under the opposite page. This part of the test was intended to explore the interaction with ancillary information as outlined in Section 4.3—"Ancillary information". Interacting with the recipe variation we provided is analogous to interacting with a variant apparatus in the context of literary criticism, while the historical background we gave for the recipe might as well be editor-provided commentary on a text. Likewise, the detachable shots of the dish could, for example, stand in for facsimiles of authentic documents in the target domain. Therefore, we argue that this is a valid method to evaluate the performance of this aspect of the design.

For the second task, we asked the subjects to open the book at a different position, where they found an incomplete recipe. More precisely, recipes for a sauce and for one ingredient were missing from this page, but accessible using digital cross references (cf. Section 4.3.2—"Supporting navigation") at the edge of the book. We told the subjects that they would need to retrieve these two recipes as well. One of the cross references took them to a location in the same book, whereas the other pointed to a recipe in a second book, so that the system prompted them to retrieve that book. This setup let us evaluate both inter- and intratextual navigation. Since the underlying motivation and execution of these cross-referencing activities are similar among our test scenario and our target domain — namely, to retrieve additional information on a specific point of a text by determining its location and then accessing it — we are convinced that this, too, is a valid method to evaluate the navigation facilities of the system.

For the third and final task, we gave the subjects the location of one more recipe in the tapas book and instructed them to gather digital versions of all recipes that were required to prepare the dish (once again, the recipe for a sauce was contained in another recipe, and had to be located using a cross

reference like in the previous subtask). To complete this task, the subjects had to use the excerption functionality of the prototype (cf. Section 4.5—"Digitization and excerption"). This part of the interaction is a direct analogy in motivation and execution to the creation of excerpts from an edition or other documents. Thus, we state that this is a valid method of measuring the performance of the system for its intended purpose.

After finishing with the tasks, we conducted semi-structured interviews directly at the table to recapitulate and discuss aspects of the experiment.

6.4 Results

We have reproduced the questionnaire in Table 6.1 along with how many participants ticked each item. The results suggest that although our participants are familiar with every-day computing technology (including touch screen devices), familiarity with large interactive surfaces and digital editions is far less common.

How often do you use a computer?
□ Seldom/never □ A few times per week □ Daily (6)

Are you familiar with multitouch-capable interfaces (e.g., *iPhone* display)?
□ Not at all □ Seen/used before (1) □ Regular user (5)

Are you familiar with digital tabletops or interactive surfaces?
□ Not at all (4) □ Seen/used before (2) □ Regular user

Are you familiar with digital editions?
□ Not at all (3) □ Seen/used before (3) □ Regular user

Table 6.1: Questionnaire of second user study

6.4.1 Ancillary information

We observed that the subjects generally had no major difficulties accessing the provided digital information in the

first task. This was also stated in the interview session. All subjects began the task by experimenting with the sliding panes, the detachable images were usually examined only after the panes had been completely explored. Roughly half of the subjects initially attempted to expand the panes by tapping the visible portion of it, whereas the other half tried dragging the tab away from the book. In some cases, their initial attempt failed due to external circumstances (e.g., the system might lose track of the touch when subjects tried to drag the pane, which caused the pane to snap back into its initial position), prompting some subjects to try out the other kind of interaction (i.e., if they had at first attempted to drag the pane, they would now tap the label, and vice versa).

Two subjects initially expressed confusion when expanding the second sliding pane with the topmost pane already expanded. In these cases, this caused the newly expanded pane to move partially off the table (due to the book's proximity to the edge of the table). Both subjects stated that they had expected the second pane to replace the first pane, not extend it. One subject stated that this expectation was also due to the fact that the contents of the second pane were slightly visible through the top pane, and that touching the second tab would just make this text fully visible in-place. (We had made the panes slightly transparent so that the user would be able to recognize when the panes obscured other digital objects on the table.)

After working with the detachable images at the opposite edge of the page, one subject came back to the expandable panes and attempted to tear one away from the book, something the prototype did now allow. He did so by fully expanding the pane, then dragging it upwards, past the boundaries of movement we imposed.

All in all, however, we conclude — and the interviews confirmed — that we found a dependable, yet not overly limiting, metaphor here.

Interaction with the detachable digital fragments was not as flawless. In most sessions, we had to prompt the subjects to explore this feature of the system. We observed — and this was confirmed in the interviews, too — that subjects did

not easily understand the rules that determined under what circumstances these fragments would show or hide a "close" button, move with the book when it was pushed around, or vanish or remain on the table when a page was turned.

Furthermore, recall that in our first implementation of the software prototype, whenever a fragment was detached from a page that "required" it, the system would immediately spawn another semi-transparent version of this fragment at the edge of the physical document to indicate that it was "still there". In the first three test sessions, this reliably caused the subjects to pause, then drag this fragment away from the book as well, thereby summoning yet another semi-transparent version if it. This cycle usually repeated a few times and left an unhappy subject with a stack of duplicated fragments. Therefore, after the first three tests, we altered this behavior so that no immediate replacement would be spawned when a fragment was detached from the book. We did not observe confusion about this aspect of the behavior after this change.

One participant expressed dissatisfaction with how easily fragments could be attached to physical documents, fearing that most of these connection would be accidental. She also pointed out that this would make it impossible to just place a document on the table at random, as one would run the risk of having all fragments in this area of the table stick to it; instead, she requested that fragments be attached to physical documents only after an explicit action on the part of the user. Considering how infrequently we can probably expect the user to actually attach fragments to physical documents, we are inclined to pursue this idea.

6.4.2 Navigation

In the second task, subjects had no problems activating the provided cross references. The message box with the referenced page number that would appear above the book for intratextual navigation was also well understood in most cases. However, the link to the recipe in another book — which summoned a document proxy object on the table, asking the user to retrieve the missing document — caused

confusion in most of the sessions. In fact, several subjects reported that they mistook the proxy object for an error message. Others would glance it over and see the displayed page number, and then proceed to open the wrong document (i.e., the document they were interacting with at this moment) at the specified page. A similar mishap happened to one participant who confused the displayed page number beneath a book (intended to give feedback that the system correctly recognized at what page the book was opened) with the page number he was supposed to navigate to.

However, despite these problems, all subjects eventually managed to retrieve the linked information and complete this task successfully.

6.4.3 Excerption

Overall, creating a digital version of an entire page worked well. Interacting with these excerpts, once created, showed the same characteristics and problems as interacting with the provided digital fragments that were introduced in the first subtask.

In one experiment, the subject had to excerpt a page to which she had navigated using a digital cross reference. When arriving at the specified page, the excerpt button overlapped with the interface item that told the location of the link target. Thus, the subject assumed that this button was part of the link UI.

6.4.4 Other observations

Three subjects attempted to begin the interaction by placing the opened book face-down on the tabletop surface. When asked, they stated that they expected the system to work like a scanner.

After this happened for the second time, we started asking if they were aware of the *StarFire* video prototype [Tognazzini, 1994] — they were not.

Occlusion was not much of an issue in this test, as the subjects would rearrange books to make room. We did not see much rearrangement of books, except for when serious

occlusion did occur (mostly with sliding panes in the first subtask), or when the focus of attention changed entirely (i.e., when turning towards the second cookbook to look up a recipe for the second subtask).

Two subjects brought up in the interviews that they would like to see some kind of explicit stand-by area for digital items on the table, i.e., a dedicated place where they could put items for later retrieval. According to them, this might be the edge of the interactive surface area.

Overall feedback was positive, especially the students from the literature department expressed great interest in the project and encouraged to pursue the path taken so far.

Chapter 7

Summary and future work

. . . in which we look back at what we have accomplished, and look forward to what remains to be explored.

7.1 Summary and contributions

In this thesis, we chronicled the design and evaluation of a system that provides a novel model of interaction for complex documents such as historical-critical editions.

Starting out from an analysis of our target audience's work processes and the status quo in the problem domain at the intersection of literary studies and HCI, we discussed the characteristics of several points in the corresponding design space. After settling on an ambitious integrative design that, as we believe, blends the advantages of physical and digital media, we took our idea through two full DIA iterations with qualitative user studies. In the process, we developed increasingly detailed prototypes.Thus, our contributions include:

- Conducting an analysis of philological work processes

from an HCI perspective.

- Developing an interaction model for existing printed documents that leverages HCI technology to make the boundary between physical and digital realms more permeable, aiming to transparently enhance established work processes with digital annotation, personalization, and effortless excerption.
- Developing a novel vision of part-digital, part-physical editions, blending the advantages of both mediums.
- Using common books as widgets for the navigation of augmenting digital information in a TUI system.

Informal feedback from members of the target user group also indicates that groups within the humanities exhibit strong interest in the future of this project, which we find very promising in itself.

7.2 Future work

This thesis has probably been just the first step of many until a productive system is completed. Many of the features we presented in Chapter 4—"Design process" still need to be implemented, most prominently those that enable the system to recognize books, page numbers, and handwritten notes. Apart from these obvious issues, however, there remain other challenges in both the near and intermediate future of this project.

7.2.1 Extensive analysis

Up to this point, the designed system has seen only qualitative evaluation, which is appropriate for the early stages of the design process. Qualitative evaluation has the advantage that it supports agile exploration of the design space better than a quantitative analysis (for which another whole set of experimental sessions may have to be conducted in

order to evaluate a change to the design), but is still suited to pinpointing design flaws. However, in order to reliably compare the system's performance to its competitors' performance, hard numbers are needed. Therefore, the design needs to undergo quantitative evaluation.

Furthermore, we also need to perform user studies that encompass a large amount of complex material as it is typical of the every-day work of our target audience. As such an experiment would require a high degree of professional experience and would probably take considerable time, we have not been able to evaluate such a scenario with our volunteer participants.

7.2.2 Integration with databases and web services

Initiatives such as *TextGrid* [TextGrid] and the *Text Encoding Initiative* [TEI] provide the infrastructure to access, encode, and exchange textual information that is suited for digital editions. We believe that the designed system's functionality could benefit from compatibility with these services.

7.2.3 Sharing

We designed the system with individual professionals in mind. Every user of the system should be able to access only the contents of the edition itself and his personal data (notes, other user-created content, information that has been automatically inferred from usage patterns and similar data); i.e., a user should not be able to access personal information that was created by another user of the system. However, under certain circumstances, it may be desirable to make such information accessible to other users. This parallels the experience we have when we borrow a book and find it heavily marked up by its owner: in such a case, we may be able to profit from this annotation by treating it as cues to what the important pieces of the text are. By virtue of the flexibility of the interaction model we proposed, such third-party annotation could be treated as just another layer of digital augmentation for the text, and be shown or hid-

den on demand. Marking up shared documents suddenly becomes an option. This also opens up theoretical usage scenarios for libraries: while it is generally undesirable to have readers scribble in printed books — thereby permanently modifying the library's property — digital markings could be effortlessly hidden or removed. Therefore, a reader could benefit from strangers' modifications to the document. Due to the digital nature of these notes, it would also be possible to share such notes not only among users of one physical copy of the document, but among all users of a specific version of the central physical document, anywhere.

7.2.4 Privacy

On the other hand, users in the aforementioned scenarios clearly need some kind of control over what information to publish and with whom. This system would probably have to be expressive enough to distinguish several use cases or addressees for shared notes, and at the same time, be lightweight enough to not interfere with casual note-taking. Group-based permission systems come to mind (e.g., the archetypical "share this with: nobody, my colleagues, everyone?" prompt).

Another problem that is present even in our system design is that of *by-catch*. For example, the system in its present form would gleefully capture the user's grocery list if it happened to be written with a digital pen on paper with an Anoto pattern. Clearly, this kind of information should never make it into the system.

Lastly, the data that is gathered from direct input and usage patterns obviously constitutes highly personal information, which has to be protected from both accidental and malevolent access through third parties. For reliable use in a production scenario, the system requires additional thought in regard to this point.

7.2.5 Copyright issues

We proposed using an overhead camera to transparently create scans of physical documents. Combined with an ability to share material with other users, there may be legal implications. The controversy over the *Google Books* project [Google Books] (in which Google scans large quantities of copyrighted books and makes them accessible online) is testimony to this problem.

7.2.6 Collaboration

In our design process, we started from the basic assumption that there would always be only one person using the system at the same time. While this is typical of the work processes we aim to enhance with this work in this thesis, it is not necessarily true in all situations. (In fact, collaborative scenarios were brought up by several interviewees.) Local and remote collaboration features could be based off the sharing facilities we outlines in the preceding sections. However, we expect collaboration scenarios to differ significantly from the single-user interaction we targeted in this thesis, and therefore did not consider it.

Appendix A

Additional documents

A.1 Paper prototype behavior

The following paragraphs describe central concepts of the initial design and its responses to a number of user actions. We used these items as a formalization in the paper prototyping sessions in addition to the initial design specification, in order to ensure consistent system behavior. Note that this only constitutes a formalization and extension of the initial design (i.e., before the paper prototyping session); we changed this behavior in subsequent iterations. Neither was it intended to be a complete specification.

Concept: digital fragment Orientation: freely rotatable, but when first displayed, oriented in such a way that the user should be able to read it (since we know/hope to know where the user sits at the table).

Concept: sliding pane Fixed at one edge of the physical document; can be dragged in one direction (perpendicular to its edge of attachment). User can show/hide it by dragging it, or toggle it by tapping it. Multiple sliding panes attached to a common edge of the same physical document stack up.

Relationship of fragments and their parents A fragment can be moved around freely in the vicinity of their

physical root object. When it is dragged out from this area, a copy is made that is no longer bound to the root object and will persist on the table even when the physical root object is removed. This is still associated with a specific page or range of pages in the physical document, so it will disappear when the parent page (or page range) is left. User can prevent this from happening by either creating a standalone copy (by dragging it out from the document's sphere of influence), or by touching it (pinning it down) with one hand when the action that would normally cause the fragment to disappear occurs. The latter action also creates a copy.

User-created notes Digitally captured pen strokes. Provide time and date of creation.

User-created excerpts Excerpts created from physical documents (see below for more info). Provides time and date of creation, document name, and page number of original passage.

Editor-provided supplementary material Facsimiles, text variants, etc. — provides document name and page number of origin. Either loosely positioned around the physical document (useful for facsimiles etc., since in general, no exact image-to-printed-text correspondence can be assumed), or rigidly positioned (e.g., commentary for specific passages of text, slide-out with labeled handles that stick out from under the book).

Concept: link Always displayed in a fixed position in relation to a physical document. Bears title of the document it refers to, and a page number. If the referenced document is known to be on the table, the display is purely informative. If the referenced document is not on the table, the system displays a small arrow next to the link text. When the user taps this active area, a proxy object is created in an empty area slightly outside the user's assumed Zone of Comfortable Reach, to serve as a reminder to retrieve the referenced document for later review. The user can create links explicitly through associating passages of text. Furthermore, all digital fragments that are separated from

their physical root object also carry a designation of origin, which also works as a back-link.

Link collapsing If several links with the same target are to be rigidly displayed alongside a document, these links are collapsed into one if the individual links would overlap or be very close (approximately 2 cm) to each other.

Concept: proxy Serves as a placeholder/reminder for a physical document that is currently unavailable. At most one proxy per physical document. If opened from a link, displays the page (or a list of pages) that were referenced from that link. If non-occluded and not in the direct vicinity of other objects, proxy appears in the position where that document was last seen. Otherwise, proxy is moved outwards from that position, towards the far end of the table, until such a position is found.

User places physical document (e.g., book) on table System recognizes new object, provides audible feedback or displays a small label on the table surface adjacent to the object to confirm correct recognition. In this step, the system is free to spawn proxies to facilitate task rehydration when no (opened) book is currently in the presumed primary working area.

User opens book System recognizes opened page, responds by displaying additional information adjacent to the book (digital fragments and supplementary material as outlined above).

User closes book System withdraws all additional information displayed on previous step.

User removes book from the table System erases corresponding digital fragments.

User places notebook on table Recently taken notes that have not yet been assigned to a single passage of text (i.e., no manually established association; up to ten most recent notes to avoid clutter) appear as sticking out from under the notebook.

User opens notebook As in the regular book case described above, the system recognizes the opened page

and the notebooks position. Alongside the user's note, the system displays time and date of capture for the existing notes, relative to today (e.g., "Two days ago"). Furthermore, for each note, the system indicates related documents and passages within these documents by displaying two groups of references alongside the note: Manually established (hard) associations, and presumed (soft) associations, based on documents that were on the table at the time when the note was taken, which documents other notes that are spatially or temporally close to the note were associated with, etc. Alongside each such note, an edge of a digital fragment appears to be sticking out from under the notebook. The user may drag on such an edge to fully reveal a digital version of the note. If such a note is released in the direct vicinity of the notebook, it slides back under it. Otherwise, upon release, it continues exists as a full digital fragment on the table, and the edge under the notebook is reestablished, to indicate that the newly created digital fragment is merely a copy.

User takes a note in notebook System captures strokes, decides whether to group it with a previous note (spatial, temporal criteria) in real-time through a clustering algorithm. If it is associated with an existing note, that note's logical area is extended.

User drags digital fragment across table Performed by touching the fragment and dragging it, without letting go of it in-between. Requires a certain small initial distance to be overcome in order to register as dragging. While dragging, the fragment resides in a layer above all other fragments on the table. Does not collide, i.e., the user can slide it across other digital fragments and under all physical objects.

User couples/decouples digital fragment to/from page User touches and drags a digital fragment on the table so that it slides under the page of an opened document from the left or right side. Upon releasing the fragment in this position, the fragment is now rigidly attached to that page (see rules for rigidly attached fragments above). User may detach fragments from physical pages in a similar fashion, by grabbing

them and dragging them straight out from under the object. When a certain initial resistance (designed to avoid accidental decoupling) is overcome, the digital fragment is promoted to a free existence on the table, and may subsequently be attached to a different position or different object altogether.

User creates an excerpt Triggered when the user places her finger in a lateral area directly adjacent to a physical document. Without leaving the surface of the table, she can then run her finger up or down alongside the document to mark a passage of text for excerption. Feedback is provided by marking the selected passage of text through a highlighted area alongside the physical object, tracing the user's finger. Can cross right over/through rigidly positioned notes.

User forges an association User performs two locally bounded, simultaneous double-taps (both taps within a half-second timeframe) both in the lateral area of physical objects described above. The system responds by displaying a newly forged link between these two passages of text, directly adjacent to the location of the respective taps, yet spacing that link far enough apart from other items in the area as to not obscure it completely.

User removes a digital fragment from the table Make "close" widget appear on all freely existing digital fragments on the table, e.g., digital representations of hand-written annotations do not show such a close button when coupled to a physical object, but exhibit one whenever floating freely on the table. Appearance/disappearance is recalculated when a drag ends.

Bibliography

Anoto. Anoto Group AB. URL http://www.anoto.com.

Maribeth Back, Jonathan Cohen, Rich Gold, Steve Harrison, and Scott Minneman. Listen reader: an electronically augmented paper-based book. In *CHI '01: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 23–29, New York, NY, USA, 2001. ACM.

Maribeth J. Back and Jonathan Cohen. Page detection using embedded tags. In *UIST '00: Proceedings of the 13th annual ACM symposium on User interface software and technology*, pages 159–160, New York, NY, USA, 2000. ACM.

George Buchanan and Tom Owen. Improving navigation interaction in digital documents. In *JCDL '08: Proceedings of the 8th ACM/IEEE-CS joint conference on Digital libraries*, pages 389–392, New York, NY, USA, 2008. ACM.

Vannevar Bush. As we may think. *Atlantic Monthly*, 176(1): 641–649, March 1945. URL http://www.theatlantic.com/doc/194507/bush.

Dino Buzzetti and Malte Rehbein. Textual fluidity and digital editions. In Milena Dobreva, editor, *Text Variety in the Witnesses of Medieval Texts*, pages 14–39, 1998.

Stuart K. Card, Lichan Hong, Jock D. Mackinlay, and Ed H. Chi. 3book: a scalable 3d virtual book. In *CHI '04: CHI '04 extended abstracts on Human factors in computing systems*, pages 1095–1098, New York, NY, USA, 2004. ACM.

Jonas Carlquist. Medieval Manuscripts, Hypertext and Reading. Visions of Digital Editions. *Lit Linguist Computing*, 19(1):105–118, 2004.

Yi-Chun Chu, David Bainbridge, Matt Jones, and Ian H. Witten. Realistic books: a bizarre homage to an obsolete medium? In *JCDL '04: Proceedings of the 4th ACM/IEEE-CS joint conference on Digital libraries*, pages 78–86, New York, NY, USA, 2004. ACM.

Andy Cockburn, Carl Gutwin, and Jason Alexander. Faster document navigation with space-filling thumbnails. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 1–10, New York, NY, USA, 2006. ACM.

Alan Cooper, Robert Reimann, and David Cronin. *About Face 3: The Essentials of Interaction Design*. Wiley, Indianapolis, IN, USA, 2007.

Katherine Everitt, Chia Shen, Kathy Ryall, and Clifton Forlines. Multispace: Enabling electronic document micromobility in table-centric, multi-device environments. In *TABLETOP '06: Proceedings of the First IEEE International Workshop on Horizontal Interactive Human-Computer Systems*, pages 27–34, Washington, DC, USA, 2006. IEEE Computer Society.

Katherine Everitt, Meredith Ringel Morris, A. J. Bernheim Brush, and Andrew D. Wilson. Docudesk: An interactive surface for creating and rehydrating many-to-many linkages among paper and digital documents. In *Tabletop*, pages 25–28, 2008.

Kenneth P. Fishkin. A taxonomy for and analysis of tangible interfaces. *Personal Ubiquitous Comput.*, 8(5):347–358, 2004.

Google Books. The Google Books project. URL `http://books.google.com`.

François Guimbretière. Paper augmented digital documents. In *UIST '03: Proceedings of the 16th annual ACM symposium on User interface software and technology*, pages 51–60, New York, NY, USA, 2003. ACM.

Jefferson Y. Han. Low-cost multi-touch sensing through frustrated total internal reflection. In *UIST '05: Proceedings of the 18th annual ACM symposium on User interface software and technology*, pages 115–118, New York, NY, USA, 2005. ACM.

Mark S. Hancock, Sheelagh Carpendale, Frederic D. Vernier, Daniel Wigdor, and Chia Shen. Rotation and translation mechanisms for tabletop interaction. In *TABLETOP '06: Proceedings of the First IEEE International Workshop on Horizontal Interactive Human-Computer Systems*, pages 79–88, Washington, DC, USA, 2006. IEEE Computer Society.

Björn Hartmann, Meredith Ringel Morris, Hrvoje Benko, and Andrew D. Wilson. Pictionaire: supporting collaborative design work by integrating physical and digital artifacts. In *CSCW '10: Proceedings of the 2010 ACM conference on Computer supported cooperative work*, pages 421–424, New York, NY, USA, 2010. ACM.

Gero Herkenrath, Thorsten Karrer, and Jan Borchers. Twend: twisting and bending as new interaction gesture in mobile devices. In *CHI '08: CHI '08 extended abstracts on Human factors in computing systems*, pages 3819–3824, New York, NY, USA, 2008. ACM.

William C. Hill, James D. Hollan, Dave Wroblewski, and Tim McCandless. Edit wear and read wear. In *CHI '92: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 3–9, New York, NY, USA, 1992. ACM.

Aldo Hoeben and Pieter Jan Stappers. Flicking through page-based documents with thumbnail sliders and electronic dog-ears. In *CHI '00: CHI '00 extended abstracts on Human factors in computing systems*, pages 191–192, New York, NY, USA, 2000. ACM.

HumTec. Brain Concept Writing at the Human Technology Centre. URL `http://www.humtec.rwth-aachen.de`.

Shaun K. Kane, Daniel Avrahami, Jacob O. Wobbrock, Beverly Harrison, Adam D. Rea, Matthai Philipose, and Anthony LaMarca. Bonfire: a nomadic system for hybrid laptop-tabletop interaction. In *UIST '09: Proceedings of the 22nd annual ACM symposium on User interface software and technology*, pages 129–138, New York, NY, USA, 2009. ACM.

Scott Klemmer, Mark W. Newman, and Raecine Sapien. The designer's outpost: a task-centered tangible interface for

web site information design. In *CHI '00: CHI '00 extended abstracts on Human factors in computing systems*, pages 333–334, New York, NY, USA, 2000. ACM.

Hideki Koike, Yoichi Sato, Yoshinori Kobayashi, Hiroaki Tobita, and Motoki Kobayashi. Interactive textbook and interactive venn diagram: natural and intuitive interfaces on augmented desk system. In *CHI '00: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 121–128, New York, NY, USA, 2000. ACM.

Chunyuan Liao, François Guimbretière, Ken Hinckley, and Jim Hollan. Papiercraft: A gesture-based command system for interactive paper. *ACM Trans. Comput.-Hum. Interact.*, 14(4):1–27, 2008.

Veronica Liesaputra and Witten H. Ian. Seeking information in realistic books: a user study. In *JCDL '08: Proceedings of the 8th ACM/IEEE-CS joint conference on Digital libraries*, pages 29–38, New York, NY, USA, 2008. ACM.

Jun Liu, David Pinelle, Samer Sallam, Sriram Subramanian, and Carl Gutwin. Tnt: improved rotation and translation on digital tables. In *GI '06: Proceedings of Graphics Interface 2006*, pages 25–32, Toronto, Ont., Canada, Canada, 2006. Canadian Information Processing Society.

Wendy E. Mackay, Guillaume Pothier, Catherine Letondal, Kaare Bøegh, and Hans Erik Sørensen. The missing link: augmenting biology laboratory notebooks. In *UIST '02: Proceedings of the 15th annual ACM symposium on User interface software and technology*, pages 41–50, New York, NY, USA, 2002. ACM.

Catherine C. Marshall and Sara Bly. Turning the page on navigation. In *JCDL '05: Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries*, pages 225–234, New York, NY, USA, 2005. ACM.

MCG. The Media Computing Group at RWTH Aachen University. URL `http://hci.rwth-aachen.de`.

M. R. Morris, A. J. B. Brush, and B. Meyers. Reading revisited: Evaluating the usability of digital display surfaces for active reading tasks. In *Proc. of TABLETOP 2007*, pages 79–86, Los Alamitos, CA, USA, 2007. IEEE Computer Society.

Tomer Moscovich and John F. Hughes. Indirect mappings of multi-touch input using one and two hands. In *CHI '08: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, pages 1275–1284, New York, NY, USA, 2008. ACM.

Christian Müller-Tomfelde, Anja Wessels, and Claudia Schremmer. Tilted tabletops: In between horizontal and vertical workspaces. In *Tabletop*, pages 49–56, 2008.

William Newman and Pierre Wellner. A desk supporting computer-based interaction with paper documents. In *CHI '92: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 587–592, New York, NY, USA, 1992. ACM.

Kenton O'Hara and Abigail Sellen. A comparison of reading paper and on-line documents. In *CHI '97: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 335–342, New York, NY, USA, 1997. ACM.

Kenton O'Hara, Alex S. Taylor, William M. Newman, and Abigail Sellen. Understanding the materiality of writing from multiple sources. *Int. J. Hum.-Comput. Stud.*, 56(3): 269–305, 2002.

Stefan Oppl and Christian Stary. Tabletop concept mapping. In *TEI '09: Proceedings of the 3rd International Conference on Tangible and Embedded Interaction*, pages 275–282, New York, NY, USA, 2009. ACM.

Anne Marie Piper and James D. Hollan. Tabletop displays for small group study: affordances of paper and digital materials. In *CHI '09: Proceedings of the 27th international conference on Human factors in computing systems*, pages 1227–1236, New York, NY, USA, 2009. ACM.

Bodo Plachta. *Editionswissenschaft: eine Einführung in Methode und Praxis der Edition neuerer Texte*. Reclam, Stuttgart, 1997.

Jun Rekimoto and Masanori Saitoh. Augmented surfaces: a spatially continuous work space for hybrid computing environments. In *CHI '99: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 378–385, New York, NY, USA, 1999. ACM.

Bill N. Schilit, Gene Golovchinsky, and Morgan N. Price. 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*, pages 249–256, New York, NY, USA, 1998. ACM Press/Addison-Wesley Publishing Co.

Chia Shen. Multi-user interface and interactions on direct-touch horizontal surfaces: Collaborative tabletop research at merl. In *TABLETOP '06: Proceedings of the First IEEE International Workshop on Horizontal Interactive Human-Computer Systems*, pages 53–54, Washington, DC, USA, 2006. IEEE Computer Society.

Jürgen Steimle. Designing pen-and-paper user interfaces for interaction with documents. In *TEI '09: Proceedings of the 3rd International Conference on Tangible and Embedded Interaction*, pages 197–204, New York, NY, USA, 2009. ACM.

Jürgen Steimle, Oliver Brdiczka, and Max Mühlhäuser. Coscribe: Using paper for collaborative annotations in lectures. In *ICALT '08: Proceedings of the 2008 Eighth IEEE International Conference on Advanced Learning Technologies*, pages 306–310, Washington, DC, USA, 2008. IEEE Computer Society.

TEI. TEI: Text Encoding Initiative. URL `http://www.tei-c.org`.

Lucia Terrenghi, David Kirk, Abigail Sellen, and Shahram Izadi. Affordances for manipulation of physical versus digital media on interactive surfaces. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 1157–1166, New York, NY, USA, 2007. ACM.

TextGrid. The TextGrid project. URL `http://www.textgrid.de`.

Bruce Tognazzini. The "starfire" video prototype project: a case history. In *CHI '94: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 99–105, New York, NY, USA, 1994. ACM.

Aaron Toney and Bruce H. Thomas. Considering reach in tangible and table top design. In *TABLETOP '06: Proceedings of the First IEEE International Workshop on Horizontal Interactive Human-Computer Systems*, pages 57–58, Washington, DC, USA, 2006. IEEE Computer Society.

John Underkoffler and Hiroshi Ishii. Urp: a luminous-tangible workbench for urban planning and design. In *CHI '99: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 386–393, New York, NY, USA, 1999. ACM.

Malte Weiss, Julie Wagner, Yvonne Jansen, Roger Jennings, Ramsin Khoshabeh, James D. Hollan, and Jan Borchers. Slap widgets: bridging the gap between virtual and physical controls on tabletops. In *CHI '09: Proceedings of the 27th international conference on Human factors in computing systems*, pages 481–490, New York, NY, USA, 2009. ACM.

D. Wigdor, G. Perm, K. Ryall, A. Esenther, and Chia Shen. Living with a tabletop: Analysis and observations of long term office use of a multi-touch table. In *Horizontal Interactive Human-Computer Systems, 2007. TABLETOP '07. Second Annual IEEE International Workshop on*, pages 60–67, Oct. 2007.

Andrew D. Wilson. Playanywhere: a compact interactive tabletop projection-vision system. In *UIST '05: Proceedings of the 18th annual ACM symposium on User interface software and technology*, pages 83–92, New York, NY, USA, 2005. ACM.

Chih-Sung Andy Wu, Susan J Robinson, and Alexandra Mazalek. Wikitui: leaving digital traces in physical books. In *ACE '07: Proceedings of the international conference on Advances in computer entertainment technology*, pages 264–265, New York, NY, USA, 2007. ACM.

Ron Yeh, Chunyuan Liao, Scott Klemmer, François Guimbretière, Brian Lee, Boyko Kakaradov, Jeannie Stamberger, and Andreas Paepcke. Butterflynet: a mobile capture and access system for field biology research. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 571–580, New York, NY, USA, 2006. ACM.

Index

Typeset April 14, 2010