TECHNISCHE UNIVERSITÄT DARMSTADT

Physical Interaction Concepts for Knowledge Work Practices

Dissertation
Zur Erlangung des akademischen Grades
eines Doktor rerum naturalium (Dr. rer. nat.)

Eingereicht von
MSC. Mohammadreza Khalilbeigi Khameneh
geb. in Teheran

Angenommen vom
Fachbereich Informatik der Technischen Universität Darmstadt

Gutachter: Prof. Dr. Max Mühlhäuser (TU Darmstadt)
Prof. Dr. James D. Hollan (UC San Diego)
Dr. Jürgen Steimle (Max Planck Institute)

Tag der Einreichung: 18.03.2014
Tag der Disputation: 08.05.2014

D17–Darmstadt 2014

Abstract

The majority of workplaces in developed countries concern knowledge work. Accordingly, the IT industry and research made great efforts for many years to support knowledge workers – and indeed, computer-based information workplaces have come of age. Nevertheless, knowledge work in the physical world has still quite a number of unique advantages, and the integration of physical and digital knowledge work leaves a lot to be desired. The present thesis aims at reducing these deficiencies; thereby, it leverages late technology trends, in particular interactive tabletops and resizable hand-held displays.

We start from the observation that knowledge workers develop highly efficient practices, skills, and dexterity of working with physical objects in the real world, whether content-unrelated (coffee mugs, stationery etc.) or content-related (books, notepads etc.). Among the latter, paper-based objects – the notorious analog information bearers – represent by far the most relevant (super-) category. We discern two kinds of practices: *collective* practices concern the arrangement of objects with respect to other objects and the desk, while *specific* practices operate on individual objects and usually alter them. The former are mainly employed for an effective management of the physical desktop workspace – e.g., everyday objects are frequently moved on tables to optimize the desk as a workplace – or an effective organization of paper-based documents on the desktop – e.g., stacking, fanning out, sorting etc. The latter concern the specific manipulation of physical objects related to the task at hand, i.e. knowledge work. Widespread assimilated practices concern not only writing on, annotating, or spatially arranging paper documents but also sophisticated manipulations – such as flipping, folding, bending, etc.

Compared to the wealth of such well-established practices in the real world, those for digital knowledge work are bound by the indirection imposed by mouse and keyboard input, where the mouse provided such a great advancement that researchers were seduced to calling its use "direct manipulation". In this light, the goal of this thesis can be rephrased as exploring novel interaction concepts for knowledge workers that i) exploit the flexible and direct manipulation potential of physical objects (as present in the real world) for more intuitive and expressive interaction with digital content, and ii) improve the integration of the physical and digital knowledge workplace. Thereby, two directions of research are pursued. Firstly, the thesis investigates the collective practices executed on the desks of knowledge workers, thereby discerning content-related (more precisely, paper-based documents) and content-unrelated object – this part is coined as table-centric approaches and leverages the

technology of interactive tabletops. Secondly, the thesis looks at specific practices executed on paper, obviously concentrating on knowledge related tasks due to the specific role of paper – this part is coined as paper-centric approaches and leverages the affordances of paper-like displays, more precisely of resizable i.e. rollable and foldable displays.

The table-centric approach leads to the challenge of blending interactive tabletop technology with the established use of physical desktop workspaces. We first conduct an exploratory user study to investigate behavioral and usage patterns of interaction with both physical and digital documents on tabletop surfaces while performing tasks such as grouping and browsing. Based on results of the study, we contribute two sets of interaction and visualization concepts – coined as PaperTop and ObjecTop – that concern specific paper based practices and collective practices, respectively. Their efficiency and effectiveness are evaluated in a series of user studies.

As mentioned, the paper-centric perspective leverages late ultra-thin resizable display technology. We contribute two sets of novel interaction concepts again – coined as FoldMe and Xpaaand – that respond to the design space of dual-sided foldable and of rollout displays, respectively. In their design, we leverage the physical act of resizing not "just" for adjusting the screen real estate but also for interactively performing operations. Initial user studies show a great potential for interaction with digital contents, i.e. for knowledge work.

Zusammenfassung

Der Großteil der Arbeitsplätze in entwickelten Ländern befasst sich mit Wissensarbeit. Die IT-Industrie und -Forschung hat entsprechend große Anstrengungen unternommen, um Wissensarbeiter zu unterstützen – tatsächlich sind computerbasierte Arbeitsplätze heute allgegenwärtig. Nichtsdestotrotz bietet die Wissensarbeit in der physischen Welt noch immer eine Vielzahl von Vorteilen gegenüber der digitalen Welt. Insbesondere die Integration der digitalen in die physische Welt lässt noch viel zu wünschen übrig. Die vorliegende Arbeit hat zum Ziel, diese Defizite unter Nutzung neuer Technologietrends, insbesondere von Tabletops und größenveränderbaren tragbaren Displays, zu beheben.

Ausgangspunkt ist dabei die Beobachtung, dass Wissensarbeiter hochgradig effiziente Praktiken, Fähigkeiten und Fertigkeiten im Umgang mit physischen Objekten in der realen Welt entwickelt haben – und zwar sowohl für wissensbezogene Hilfsmittel (Bücher, Notizblöcke, ...) als auch für nicht wissensbezogene (Kaffeetassen, Büromaterialien, ...). Unter den wissensbezogenen Objekten stellen papierbasierte Objekte die bei weitem relevanteste (Über-)Kategorie dar – Papier ist der klassische analoge Informationsträger. Wir unterscheiden hier zwei Arten von Praktiken: *Kollektive* Praktiken betreffen die Anordnung von Objekten bezogen auf andere Objekte auf dem Tisch, während *spezifische* Praktiken auf einzelnen Objekten operieren und diese gewöhnlich verändern. Die erste Kategorie wird hauptsächlich für ein effektives Management des physischen Arbeitsplatzes eingesetzt – z.B. werden gewöhnliche Objekte häufig auf dem Schreibtisch verschoben, um den Arbeitsplatz zu optimieren - oder zur effektiven Organisation papierbasierter Dokumente auf dem Tisch – z.B. stapeln, auffächern, sortieren, ... Die zweite Kategorie befasst sich mit der spezifischen Manipulation von physischen Objekten in Bezug auf die aktuelle Aufgabe, z.B. Wissensarbeit. Weitläufig genutzte Praktiken beinhalten nicht nur das Schreiben auf, Annotieren von oder räumliche Anordnen von Papierdokumenten, sondern auch die Manipulation derselben, zum Beispiel durch drehen, falten, biegen, usw.

Verglichen mit der Vielzahl solcher etablierter Praktiken in der realen Welt sind die Praktiken in der digitalen Welt heute noch dominiert durch die Eingabe mit Maus und Tastatur. Die Maus war zwar ein so großer Fortschritt, dass Forscher verleitet wurden, ihre Verwendung "Direktmanipulation" zu nennen, doch aus heutiger Sicht ist klar, dass sie immer noch eine sehr hinderliche Indirektion darstellt. In diesem Licht kann das Ziel dieser Arbeit umformuliert werden als: Erforschung neuer Interaktionskonzepte für Wissensarbeiter, die i) das Potential der flexiblen und di-

rekten Manipulation physischer Objekte für eine intuitivere und ausdrucksstarke Interaktion mit digitalen Inhalten (sehr ähnlich derjenigen in der realen Welt) ausnutzen, und ii) die Integration digitaler und physischer Arbeitsplätze verbessern. Hierbei werden zwei Forschungsrichtungen verfolgt: Erstens werden die *kollektiven* Praktiken von Wissensarbeitern am Schreibtisch untersucht; dabei wird unterschieden zwischen wissensbezogenen (Papier / Dokumente) Objekten und nicht wissensbezogenen Hilfsmitteln – letzteres wird im Folgenden "table-centric approaches" genannt, Basis sind interaktive Tabletops. Zweitens untersucht die Arbeit *spezifische* Praktiken für die Arbeit auf Papier, sie konzentriert sich dabei auf wissensbezogene Aufgaben aufgrund der spezifischen Rolle von Papier in diesem Bereich - dieser Teil hat den Titel "paper-centric approaches". Dabei werden die so genannten Affordances von papierähnlichen Displays ausgenutzt, genauer gesagt von größenveränderbaren, z.B. roll- und faltbaren, Displays

Der "table-centric approach" führt zu der Herausforderung, interaktive Tabletops mit der allgemein etablierten Nutzung konventioneller (Tisch-)Arbeitsplätze zu kombinieren. Wir führen zunächst eine exploratorische Nutzerstudie durch, um die Verhaltens- und Nutzungsmuster der Interaktion mit sowohl physischen als auch digitalen Dokumenten auf Tabletops bei der Ausführung von Überblicks- und Gruppieraufgaben zu untersuchen. Basierend auf den Ergebnissen der Studie entwickeln wir zwei Interaktions- und Visualisierungskonzepte – PaperTop und ObjecTop – die die spezifischen Praktiken papierbasierter bzw. kollektiver Arbeit berücksichtigen. Die Effizienz und Effektivität werden jeweils in einer Reihe von Benutzerstudien untersucht.

Wie erwähnt nutzt der "paper-centric approach" ultradünne größenveränderbare Displays. Wir tragen hier wiederum zwei neue Interaktionskonzepte bei – genannt FoldMe und Xpaaand - - welche unsere Antwort auf den Design-Space von zweiseitigen falt- bzw. rollbaren Displays darstellen. Im Design nutzen wir den Vorgang der Größenveränderung nicht nur, um die Bildschirmgröße anzupassen, sondern auch um interaktiv Operationen durchzuführen. Erste Nutzerstudien zeigen ein großes Potential für die Interaktion mit digitalem Inhalt, z.B. für die Wissensarbeit.

Acknowledgments

My research in the last five years that resulted in this manuscript would have been impossible without the support of my adviser, colleagues, students, and family. I am very grateful to all of them and those who contributed to this thesis, in one way or another.

First, I would like to thank my supervisor Max Mühlhäuser who made it possible to even start this endeavor and for his unlimited support, guidance, and encouragement throughout the years of my research. I am also very grateful to James D. Hollan for acting as co-referee and his constructive and supportive feedback that informed my overall research.

Jürgen Steimle as my third advisor has played equally important role, co-supervising my research and challenging me to always give my best. Without his ideas and enthusiasm I would not be the same researcher today. Thank you!

I was very fortunate to work in a friendly, fruitful, and fun environment, i.e. Telecooperation (TK) Lab at Technische Universität Darmstadt, in which I collaborate with an excellent group of colleagues: Immanuel Schweizer, Kamill Panitzek, Dirk Bradler, Simon Olberding, Thorsten Strufe, Erwin Aitenbichler, Daniel Schreiber, Felix Heinrichs, Elke Halla, and Nina Jäger. I want to thank all other former and current TK members that made my time there such a wonderful experience. Likewise, I sincerely thank all the students I worked with for their valuable creative contributions: Tarik Tahiri, Christoph Bauer, Patrik Schmittat, Philpp Sowinski, Wolfgang Kleine, Manolis Pavlakis, and Sebastian Günther.

I am most grateful to my awesome fellow research group members and–most importantly–friends: Jochen Huber, Roman Lissermann, Florian Müller, Jan Riemann, and Martin Schmitz.

I would like to thank family for all their remote support. Thank you to my parents, brother, and sister for all their encouragement and kind help when I needed it. I wish to extend special thanks to the Payandeh family for their continuous support.

And above all, I want to thank my wonderful wife Niloo who has been with me from the beginning of this endeavor. She has given me support when I needed it most, distraction when I needed to get my head off work and love – all the time. Last but not least, a big thank goes to my little daughter, Lena, that made nights of writing this thesis an adventure.

Contents

Chapter 1

Introduction

Contents

The present thesis explores novel concepts for interaction with digital (computer-based) content. These concepts exploit the flexible manipulation of physical objects in the real space. They are designed for and integrated into the natural ecology of knowledge workers.

This chapter introduces the reader to the motivation, challenges and approaches that set the foundations of this thesis. It then presents the contributions and the structure of the thesis.

1.1 Motivation

Since the advent of direct manipulative interfaces [Shneiderman 1982], knowledge work settings typically consist of two quite distinct worlds: digital objects confined in desktop PCs, and the physical world of everyday objects spatially distributed on tables. The digital objects, mainly documents and folders, are graphically represented on the computer monitor in a way that resembles a physical desktop. All interactions with digital objects are channeled through indirect input devices – such as mouse and keyboard – which enable the manipulation of Graphical User Interfaces (GUIs) in the WIMP (Windows, Icons, Menu, Pointer) paradigm. While successful and well-established, these input devices do not incorporate the highly developed manipulation skills, practices and dexterity of their users in the physical world. For instance, the mouse allows for controlling only a single object at any given moment in time, making it a generic, time-multiplex input device [Fitzmaurice 1995].

Moreover, feedback related to mouse or keyboard input is mostly provided as visual information that is displayed separately on the computer monitor.

Over the past two decades, HCI researchers have sought to develop interfaces that improve over the WIMP paradigm. Development of these so-called, post-WIMP interfaces [Jacob 2008] has coincided with advances in display and input technologies.

Chief among the Post-WIMP enabling technologies is the advent of multitouch input that can be considered as a breakthrough in the last decades. It can be considered as one step closer to direct manipulation by allowing interaction with digital contents through bare fingers and hands without the indirection introduces by a mouse or keyboard. Given the advantages of the touch input, it has been implemented on a various devices with different form factors. Despite rapture of multi-touch interaction, it has also contributed to impoverishing the tangible physicality of interfaces. Multitouch interaction or *Picture Under Glass* interaction paradigm [Victor 2011], as it has been described, scarifies all the tactile richness of working with the fingers.

Tangible User Interfaces (TUIs) are another example of post-WIMP interfaces. They aim to mitigate the problem of lacking tangibility by coupling the bits with everyday physical objects [Ishii 1997]. TUIs refer to interfaces that enable controlling and sometimes displaying digital information using physical objects. Such interfaces encourage bimanual interactions and offer space-multiplex input [Fitzmaurice 1995]: Different parts of the body (mainly the two hands) are used to concurrently interact with different objects that represent dedicated functions, similar to how we interact in the real world. Although direct touch input also provides for space-multiplex input, TUIs additionally take advantage of natural affordances offered by physical objects as well as the embodiment available through tangibility [Hornecker 2006].

The goal of this thesis is to explore tangible interaction concepts based on flexible manipulation of physical real world objects. Leveraging well-established practices of working with physical objects, these concepts support knowledge work activities spanning across the digital and the real worlds and thus, aiming to dissolve the boundaries between two worlds by enabling more physical ways of interacting with computers.

1.1.1 Knowledge Work Practices

The physical desktop workspace of knowledge workers (KWs) usually comprises of a wide variety of physical objects that are placed on the desk. A substantial part of these objects is content-related – i.e., media containing the knowledge worked on – and we often witness it as a myriad of paper-based artifacts such as, printed documents, books, and magazines. For the present thesis, it is equally important to respect the second category of physical objects that usually populate a desk: the content-unrelated ones – such as office stationery, coffee cups, plants, and hardware components of desktop computers. Both categories of physical objects are spatially arranged across the table, more or less intuitively. In case of paper-based artifacts, a very common form of organization is piling them on each other [Malone 1983]. Such a spatial arrangement of objects usually results in a *muddle*, where the knowledge work takes place. The tasks demanded of knowledge workers are rich and highly varying [Kidd 1994]. To undertake such tasks, knowledge workers continually alter the arrangement of their workspace. For instance, the physical objects are frequently rearranged and shuffled on tables.

A typical task assigned to knowledge workers can be considered to consist of three main phases. First, they have to retrieve a number of content-related objects (mostly paper-based documents) from various resources – such as piles and then bring them to the working area of their table. The second phase is while the task is processed. In this phase, KWs deeply engage in documents and employ more fine-grained methods for various purposes – such as navigation between and within documents through flipping or folding techniques [Marshall 2005a]. After processing the task, the results need to be forwarded to another KW or sorted and organized back to their corresponding piles or folders. In all these phases, plenty of content-related and -unrelated physical objects are used together. For example, a knowledge worker interacts with paper documents while drinking a coffee or using office stationery – such as pens and staplers.

In general, the practices employed in these three phases can be divided into two categories: The first category encompasses those practices that are performed on the table level, namely, laying out, shuffling, piling and grouping physical materials on the desktop surface. Knowledge workers thereby exploit the *physical space* available on their tables for these practices [Kirsh 1995]. Research has shown that being able to spatially lay out physical materials – in particular, printed documents – is one of the salient advantages of physical workspaces compared to digital one [O'Hara 1997, Kidd 1994]. Spatial layout of materials plays an important role in knowledge work as it provides a rich mental picture or an overview of the workspace as well as strong and immediate contextual cues to rehydrate task state [Kidd 1994]. Moreover, spatial layout is also *highly dynamic*. Knowledge workers continually move materials particularly printed documents in and out of the center of their attention or temporarily create and manipulate piles placed on more peripheral region of their workspace.

The second category of practices is based on the flexible manipulation of paper. Among materials constituting the physical desktop workspace, paper is one of the primary and most widely used materials. Despite the prediction of its early demise, paper remains ubiquitous and widespread in knowledge work settings as it is documented in the seminal work by Sellen and Harper [Sellen 2003]. Even at the beginning of the 21st century, knowledge work is still tightly coupled with paper [Steimle 2012a]. This is because of the unique inherent affordances of paper over digital technologies. Paper is light, thin and malleable and thus knowledge workers can easily adapt its shape to suit the situation that is comfortable for reading. Paper is tangible and physically embodies information. Knowledge workers leverage the tangibility of paper while manipulating it. For example, during active reading, they utilize established practices – such as flipping, folding, thumbing or bending – to fluidly navigate within a document [Marshall 2005a]. Moreover, paper can be flexibly rearranged in the physical space, allowing for easy navigation between multiple documents or easily interweaving annotating while reading.

These practices, dexterity, and skills of working with physical objects have evolved over many years and became well-established. Paper in particular, but also content-unrelated objects, offer rich physical affordances that enable an easy and intuitive manipulation. Compared to these efficient practices and skills, the way KWs interact with the digital world confined in desktop PCs is, in many respects, a step backward. All interactions with digital content are channeled through indirect input devices (e.g., a mouse), and feedback (output) is displayed separated from the input on the computer monitor. Therefore, it is of great value for knowledge workers to extend digital computation and interaction to the physical space of everyday life. The present thesis is an effort towards connecting both worlds so that KWs can use the well-established practices for intuitive and direct interaction with physical object in the context of digital content. Throughout this thesis, we employ table and paper-centric practices as foundations for designing interactions to facilitate knowledge work activities. Our designs are motivated by the novel interaction paradigms offered by advances in display and input technologies.

1.1.2 Interaction Paradigm Shift

Display and input technologies have advanced quite rapidly in recent years. This led to new display form factors – e.g., mobile devices, digital tablets, interactive tabletops, and wall-size displays – and novel interaction styles, in particular to the success of multitouch devices. Given the well-established practices of knowledge workers that are mainly centered around tables and paper, we are especially interested in addressing challenges stemming from the integration of two emerging classes of displays – namely, large horizontal interactive displays (tabletops) and paper-like flexible displays – into office workspaces. Both are discussed below in detail.

Tabletop Computing

Given the ubiquity of tables in office and home environments, researchers have long been investigating the integration of computers into tabletops. As highlighted by Müller-Tomfelde and Fjeld [Muller-Tomfelde 2012], in the last two decades, tabletop computers have improved at a rapid pace and reached a high level of maturity in terms of user-related and technical issues. Such computers are capable of displaying a digital world that closely resembles the appearance and look of physical objects. They also allow for manipulation of objects using direct (multi)touch input. However, as mentioned before, the directness offered by multitouch input does not go beyond touching and swiping the display and is unable to "reach" the actual objects.

As a consequence and also due to the horizontal form factor of tabletops, considerable research has focused on integrating aspects of physicality into tabletop interfaces. Thus, many tabletop systems incorporate the use of tangible elements in design of their interfaces, resulting in so-called *hybrid surfaces* [Kirk 2009]. In such interfaces, the tangible elements are themselves part of the system and tightly coupled to the digital objects represented on tabletop surface. While prior hybrid tabletop systems [Wellner 1993, Ullmer 1997, Underkoffler 1999, Jordà 2007] provided first insights, there is still exists a lack of understanding on how interactive tabletops can be effectively integrated into desktop workplaces. More explicitly, concurrent use of the content-related and content-unrelated physical objects in combination with digital objects on interactive tabletops gives a rise to a myriad of interaction design challenges that are not yet fully explored. Therefore, one research direction of this thesis considers the integration of tabletop computers into office workplaces. This research direction particularly addresses the following research questions:

1) How can well-established knowledge work practices of working with ordinary everyday objects (particularly paper-based media) be effectively transferred to hybrid workspaces augmented with digital tabletops?

2) How to deal with hybrid tabletop workspace settings which their surface are

populated with physical objects? In particular, how to deal with occlusion of screen contents on tabletops by physical objects?

Successful integration of interactive surfaces into everyday life depends on addressing and providing practical solutions to these questions.

Paper-like Computing

As mentioned above, paper has unique physical properties over electronic platforms that explain its ubiquitous usage. Being thin and light in weight, foldable as well as easy to use, share and lay out in space are the most prominent ones identified in various studies [Sellen 2003, O'Hara 1997, Bondarenko 2005, Marshall 2005a]. These key affordances have motivated a large body of research projects exploring the role of paper documents as interfaces to the digital world (for an extensive review, see the work by Steimle [Steimle 2012a]). Such interfaces, typically utilize paper as documents with some form of digital enhancements to bridge the gap between printed and digital contents.

Given the high degree of tangibility and flexibility offered by paper, researchers have taken a step beyond using paper as a medium for reading and writing. Considerable research has focused on using passive paper-based materials – such as paper boards, sheets, cards – as interfaces to interact with computers. In such interfaces, affordances of paper permit intuitive, direct and bimanual interaction with digital contents. The recent advent of electrophoretic ink (E-Ink) and light-emitting polymer technologies has sparked a new surge of interest in investigating active paper-like displays in this vein of research. These recent technological advances allow for creating displays that offer many potential advantages such as ultra-thin profiles and lightweight yet robust. More importantly, they provide the ability to bend, flex, conform, roll, and even fold a display not only for ease of mobility but also as a novel input modality.

The ever growing trend of creating ultra-thin display technology promises a future, where fully flexible digital paper will be widely available in desktop workspaces. This allows KWs to use natural practices and skills of working with paper – such as folding, bending, leafing – to interact with digital contents visualized on flexible displays or devices featuring such displays. Therefore, the second research direction followed in this thesis is to explore interaction concepts based on the way we physically manipulate paper. In particular we address:

1) How can natural practices of working with paper such as folding and rolling be used as new modes of interaction?

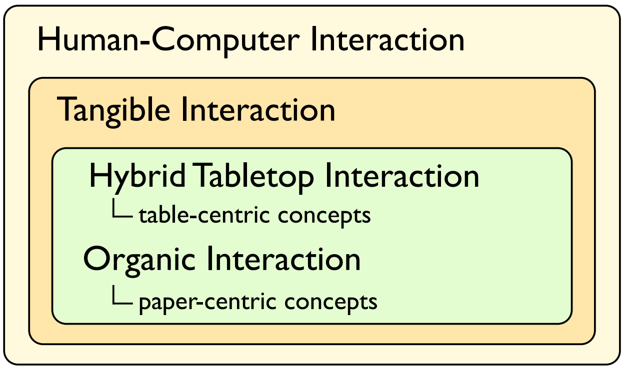

Figure 1.1: Research Context

2) How should deformation of future foldable and rollable displays be tailored for different application contexts? How can (re)size and (re)shape be effectively incorporated into interface design?

These advances in both tabletop and paper-like computing areas can potentially bring displaying digital information and interaction possibilities to the very artifacts of our physical space such as tables and paper. Given ubiquity and salient role of these two artifacts in knowledge work settings, this thesis aims to effectively integrate the digital world of computing into the physical world of surface and paper and thus, improving not only the efficiency but also expressiveness and user experience of interaction. The present thesis is an effort toward a desktop workspace of future, where these novel display technologies are commonplace: the physical desktop is augmented with an interactive tablecloth, on which not only physical but also digital objects as well as a number of paper-like displays can be effectively manipulated.

1.2 Research Context, Method, and Contributions

As mentioned above, this thesis explores and analyzes novel interaction concepts based on flexible manipulation of physical objects to support knowledge work activities. In terms of the general scientific disciplines, it primarily contributes to the field of Human-Computer Interaction, the study of interaction between people and computers and the design of novel interface approaches. More explicitly, the focus of this thesis is situated in the field of Tangible User Interfaces [Ishii 1997] and Tangible Interaction [Hornecker 2006], where digital information and manipulation take physical representation and forms. Within this area, the two sets of research

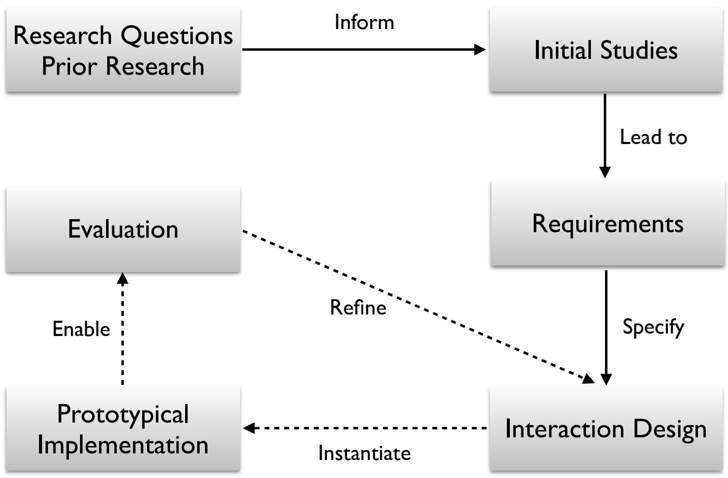

Figure 1.2: Research Methodology

questions presented above contribute to the fields of Hybrid Tabletop Interaction (i.e., use of tangible controls for surface-based interaction [Kirk 2009]) and Organic Interaction (i.e., use of non-planar displays as a primary means of output as well as input [Holman 2008]). Figure 1.1 illustrates the research context of this thesis.

The main contribution of this thesis is the novel interaction concepts designed to improve both efficiency and user experience of knowledge work activities spanning across both realms. This is achieved by studying digital augmentation of tables and paper as the most ubiquitous and crucial components of knowledge work activities. Contributions provided in response to the aforementioned research questions are empirically grounded by exploring the natural behavior of users and examining literature in the respective field. Then, based on empirical qualitative and quantitative findings, we iteratively design and analyze interaction concepts with potential users (i.e., user-centered design approach [Norman 1986]). More precisely, each design iteration started with an initial study. This led to a set of requirements. We then designed novel interaction concepts that were prototypically implemented. Finally, the concepts were evaluated with potential users in order to test the validity of the assumptions. In the evaluations, we opted for both quantitative and qualitative data collection methods to measure both efficiency and user experience. A schematic view of the research method followed in this thesis is depicted in figure 1.2. In this way, this thesis provides the following contributions in detail.

A. Table-centric Concepts

In this research direction, we investigate how digital tabletops can seamlessly and effectively integrate physical desktop workspaces of knowledge workers. Since knowledge work includes concurrent interaction with physical and digital media, we begin our investigation by analyzing such combined media usage patterns on interactive tabletops. More explicitly, we contribute a thorough user study in which the affordances and tradeoffs involved in how users spatially group and arrange items are examined. Moreover, we investigate the impact of occlusion of screen contents by physical media and analyze users' strategies while dealing with it. The results of this study serve as foundations for design of two other contributions of this research direction namely: *PaperTop* i.e, a hybrid piling interface for tabletop system and *ObjecTop* i.e. an occlusion-aware user interface for hybrid tabletop systems. Both interfaces are evaluated in terms of efficiency and user experience in a series of user studies. The evaluation show that ObjecTop is faster for finding objects, decreases interaction with physical objects when resolving problems of occlusion, requires less effort, and is less frustrating.

B. Paper-centric Concepts

This research direction considers the integration of the digital world of computing with the physical world of paper. One route toward this goal is based on flexible flat panel technologies [Crawford 2005] that aim to develop displays that resemble the superior handling and flexibility of real paper. We contribute two flexible displays interfaces that are based on folding and rolling metaphors namely, *FoldMe*: a foldable display concept and *Xpaaand*: a rollable display concept. With analysis of their design space, we show how users can physically interact with foldable and rollable displays (or devices featuring such displays) and depict plausible combinations of such deformations with direct touch input. Moreover, we design a set of interaction principles and concepts to accomplish basic tasks. Our display concepts are evaluated in exploratory user studies in which we test our concepts with functional prototypes. The user studies show that physically resizing the display through folding and rolling offer a high degree of expressiveness and enjoyment while interaction.

1.3 Thesis Outline and Publications

This thesis is structured in three main chapters. **Chapter 2 and 3** focus on the integration of interactive tabletops with the physical desktop and accordingly, present our contributions mentioned above (section 1.2 A). More precisely, **Chapter 2** includes our concepts about the integration of physical and digital media on interactive tabletops. We present a user study investigating concurrent use of both media types on tabletops and drive a set of requirements for hybrid tabletop systems. Moreover, in this chapter we introduce the PaperTop interface for hybrid document piling on tabletops.

In **Chapter 3**, we describe ObjecTop, a collection of occlusion-aware techniques to support users while interacting with digital and physical objects on tabletop display surfaces. These techniques provide awareness, access and interaction with occluded digital objects. We report an evaluation of ObjecTop compared with a conventional interface in which we discuss its advantages and disadvantages over an unaided tabletop system.

In **Chapter 4**, we explore physical interaction concepts with flexible displays and present our contributions stated in section 1.2 B. More explicitly, we present two flexible display concepts based on folding and rolling metaphor. We depict a physical design space of such displays together with a number of interaction and application concepts. At the end of this chapter, two exploratory evaluations are presented in which we examine our display concepts. **Chapter 5** includes conclusions and possible future research directions.

Contents, ideas and figures presented in the three main chapters have published previously in proceedings of international conferences such as, ACM SIGCHI Conference on Human Factors in Computing Systems (CHI), ACM Conference on Interactive Tabletops and Surfaces (ITS), ACM Conference on Tangible and Embedded and Embodied Interaction (TEI) as well as international and national workshops and scientific magazines.

Contents presented in chapter 2 are published in [Khalilbeigi 2010, Khalilbeigi 2009]. The results of the user study on hybrid media usage on interactive tabletops are published in [Steimle 2010b, Steimle 2010a]. Parts of chapter focusing on physical occlusion are published in [Khalilbeigi 2012b, Khalilbeigi 2013]. Paper-centric interaction concepts are published in [Khalilbeigi 2012a, Khalilbeigi 2011, Steimle 2012b]. Figure 1.3 provides a written and visual summary of which materials are used in the thesis chapters along with the selected corresponding publications.

Thesis Chapters | Corresponding Publications

1. Introduction

Table-centric Concepts

2. Integrating Physical-Digital Media on Tabletops

Steimle, **Khalilbeigi,** Mühlhäuser and Hollan. Physical and Digital Media Usage Patterns on Interactive Tabletop Surfaces. In Proc. ITS 2010.

Khalilbeigi, Steimle and Mühlhäuser. Interaction Techniques for Hybrid Piles of Documents on Interactive Tabletops. In Proc. CHIEA 2010.

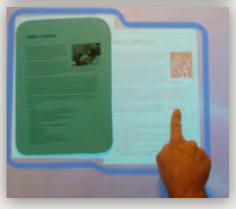

Steimle, **Khalilbeigi** and Mühlhäuser. Hybrid Groups of Printed and Digital Documents on Tabletops: A Study. In Proc. CHIEA 2010.

Khalilbeigi, Steimle and Mühlhäuser. Interaction Support for Hybrid Groups of Paper & Digital Documents on Tabletops. In Proc ITS 2009 (adjacent CD).

3. Managing Physical Occlusion on Tabletops

Khalilbeigi, Schmittat, Mühlhäuser and Steimle. Occlusion-aware Interaction Techniques for Tabletop Systems. In Proc. CHIEA 2012.

Khalilbeigi, Steimle, Riemann, Dezfuli, Mühlhäuser and Hollan. ObjecTop: Occlusion Awareness of Physical Objects on Interactive Tabletops. In Proc ITS 2013.

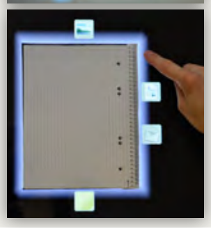

Paper-centric Concepts

4. Exploring Physical Interactions for Paper-like Displays

Khalilbeigi, Lissermann, Mühlhäuser and Steimle. Xpaaand: Interaction Techniques for Rollable Displays. In Proc. CHI 2011.

Khalilbeigi, Lissermann, Kleine and Steimle. FoldMe: Interacting with Double-sided Foldable Displays. In Proc. TEI 2012.

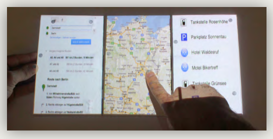

Steimle, Lissermann, Olberding, **Khalilbeigi**, Kleine and Mühlhäuser. Be-greifbare Interaktionen mit grösenveränderbaren Bildschirmen. i-com Article Oldenbourg 2012.

5. Conclusions

Figure 1.3: Thesis Outline and Corresponding Publications

Chapter 2

Integrating Physical-Digital Media on Tabletops

Contents

In this chapter, we will explore novel, more usable and appealing interaction concepts for knowledge work from a *table-centric* perspective. Given the recent technological advances in creating interactive surfaces, digital tabletops are becoming common components of knowledge work activities. Despite these computing advances, use of the physical media in knowledge work, as documented by Sellen and Harper [Sellen 2003], remains pervasive. Physical media or more precisely the paper medium, provides key affordances that are not supported by current computing systems. Among others, rich physical embodiment of information, support for easy and intuitive bimanual navigation and spatial organization, and enabling intuitive annotation are examples of the unique affordances of paper.

As a consequence, physical media is often involved concurrently with digital media. As an example, while reading a printed document users increasingly access digital information. Because of the ubiquity of its occurrence, we are particularly interested in use cases in which physical and digital medias are both involved in one activity. Although introducing digital tabletops into the knowledge work setting allows for a much greater spatial proximity between digital and physical realms as well as displaying and arranging digital items much like physical ones, it gives rise to a myriad interaction design challenges. One primary yet important question is how users spatially arrange and group both digital and physical items. Moreover, it is not clear how well-established practices of working with paper, such as piling or browsing, will be reflected while working with a hybrid set of documents on tabletops.

To this end, the present chapter addresses these questions and examines the affordances of such a combined used of physical and digital media. Through an in-depth user study, we identify a set of challenges to be addressed in novel user interface concepts that support hybrid media usage on tabletops. We then propose novel interface concepts that leverage the ways we physically manipulate printed documents for controlling digital ones. The concepts are coherently implemented in *PaperTop*, a system that facilitates hybrid piling on interactive tabletops. In summary, the main contributions of this chapter are

- exploration of physical and digital media usage patterns on interactive tabletops through an in-depth user study
- design of novel hybrid interaction and visualization concepts to support manipulation of hybrid media piles
- integration of these concepts into an overall approach, called *PaperTop*, capable of tracking multiple sheets of paper on, around, and above the tabletop surface

- evaluation of these techniques through an initial user-feedback session

The reminder of this chapter is structured as follows. We begin by reviewing related work in section 2.1 pertaining to the main contributions presented in this chapter. Particularly, regarding the exploratory study, we discuss the previous work that investigates the affordances of physical and digital media. We then discuss desktop as well as tabletop interfaces aimed at bridging the digital-physical gap.

Based on related work, in section 2.2 we present our exploratory user study to systematically understand how people work in a hybrid setting. We are particularly interested in examining the affordances and trade-offs involved in how users spatially arrange and group both types of media. Among others, the study shows that users intentionally create occlusion in the form of piles in order to express semantic relationships in spatial arrangements. *Hybrid piles* of printed and digital documents were the most dominant and frequent form of arrangement. It was also found that manipulation of hybrid piles (such as browsing, moving, getting overview) is a tedious task. Based on these and other findings, we set out several design requirements.

In section 2.3, we present our interaction concept supporting hybrid piles of digital-physical documents on interactive surfaces. More concretely, we first discuss design challenges and then present our visualization and interaction techniques to support manipulation of hybrid piles. Our techniques include a bubble visualization technique that alleviates piling activities such as creation, flexible reorganization, and fluid transitions between various pile forms based on the proximity and spatial arrangement of documents.

Details of our prototypical implementation of the techniques and algorithms to track multiple sheets of paper in the 3D space above the tabletop are described in section 2.4. In section 2.5, we report on an early user-feedback session evaluating our techniques with a number of experts. We close this chapter by a conclusion presented in 2.6.

Contribution Statement: Most of the work discussed in this chapter is based on [Khalilbeigi 2010, Khalilbeigi 2009, Steimle 2010b, Steimle 2010a]. I have led the design, implementation, and evaluation of PaperTop, published in [Khalilbeigi 2010, Khalilbeigi 2009]. The study presented in 2.2 has been initiated and led by Jürgen Steimle and I, as the second coauthor of the corresponding publications [Steimle 2010b] [Steimle 2010a] have significantly contributed in design, conducting, and analysis of its results. As with any other scientific publication, all of the other authors, Max Mühlhäuser and James D. Hollan, have contributed significantly.

2.1 Related Work

To place our work in context, we analyze three main areas of the literature. We begin by discussing prior studies investigating affordances of physical and digital media. Drawing upon findings from these studies, we formulate several research questions to be addressed by the exploratory study presented in 2.2. We continue by revisiting two areas of previous work related to the design of the PaperTop interface. The first area includes studies that assess how people organize their information in real workspaces as well as prior desktop computing systems that assist managing digital information collections. While not directly relevant, this vein of research provides a set of guidelines that influenced the design of the PaperTop interface. Second area incorporates previous tabletop systems supporting hybrid physical-digital settings. We analyze the strengths and weaknesses of each and derive a set of design guidelines.

2.1.1 Affordances of Physical and Digital Media

Understanding the natural interaction practices of knowledge workers and physical media usage patterns have been largely discussed in the literature [Sellen 2003, Sellen 1997, Whittaker 2001, O'Hara 1997]. The seminal work by Sellen and Harper on the myth of the paperless office [Sellen 2003] revealed the inherent characteristics as well as the well-established practices of using paper. They investigated and systematized the unique affordances of the paper medium. Chief among them are the rich physical embodiment of information, support for easy and intuitive bimanual navigation and organization, facilitation of communication and collaboration, enabling intuitive annotation using pens, and providing high resolution and contrast for easy reading. The authors argued that the usage of paper remains pervasive in knowledge workspaces, and rather than pursuing the idea of paperless offices, digital tools and novel technology should work *in concert with paper media* in order to make optimal use of both media types.

O'Hara and Sellen [O'Hara 1997] compared reading from paper and reading online through a laboratory user study with ten office workers. They identified key advantages of reading on paper such as support for effortless annotation while reading, quick navigation, and flexibility of the spatial layout. Their findings imply that digital systems should support more flexibility and control in spatial layouts.

With the advent of interactive surfaces, researchers employed similar approaches to investigate fundamental aspects and look at the ways people organize digital media and work around digital surfaces. Scott et al. [Scott 2004] performed an in-depth analysis of territoriality in collaborative tabletop workspaces. They found personal territories to be different from shared territories. Areas for storing items are integrated within these territories and change their positions over time. Other studies analyzed the impact of table size [Ryall 2004] and orientation of media [Kruger 2003]. Although informative, these studies have only focused on the way people manipulate virtual objects on tabletops.

Probably the previous work most related to our study, presented in this section, are those that assessed affordances of physical versus digital media on interactive surfaces [Terrenghi 2007, Piper 2009]. Terrenghi et al. compared the affordances of interacting with digital and physical media on tabletops [Terrenghi 2007]. Their study investigated the ways manipulation of physical versus digital media are fundamentally different from each other. Participants carried out a puzzle task and a photo sorting task in two different modes: with physical items in three dimensional space and with virtual items on interactive tabletops. Analysis of basic tasks with photos revealed distinctive differences in the ways people interact with digital and physical photos. Most importantly, they noted that one-handed interaction predominated with digital media, even though multitouch interaction is specially designed to support bimanual interactions. This might be partly explained by the lack of haptic feedback and by the restrictions of a two-dimensional surface. Another important finding is that due to the lack of physicality in the digital condition, participants performed more effortful strategies to assess the properties of objects for instance, the quantity of objects in a pile. They observed that participants needed to spread out piles to visually judge the quantity and content of the objects they were about to deal with. The authors argued that such fundamental differences between both realms call for the design of novel ways to use physical affordances as a design resource.

Piper and Hollan [Piper 2009] compared the affordances of digital educational material on tabletops with traditional paper handouts in collaborative study practices. The study was conducted with ten pairs of undergraduate students in two conditions: a paper condition (in which students used a set of paper diagrams) and a digital condition (in which students worked with digital images on a tabletop). One surprising finding was that students in the digital condition incorporated their personal study materials, such as paper notes and notebooks, together with the digital materials on the tabletop. They also found that the more ephemeral character of digital ink annotations encouraged students to spontaneously create drawings to support discussion, more than with paper. Authors argued that both digital and paper medias have unique and complementary affordances for small-group studies, which one needs to take into account while augmenting and supporting educational ecologies. In summary, our examination of literature in this area of research reveals that paper media

- provides a rich physical embodiment of information,
- affords a strong haptic feedback,
- supports intuitive and flexible bimanual navigation and organization,
- allows for spatially laying out information.

It was also found that digital systems aiming to support paper-based practices should work in concert with paper media rather than following the paper-less office vision. The novel interaction paradigm offered by interactive tabletops incorporate some of the affordances of paper in the digital world. However, due to the lack of physicality, the interaction with digital 2D objects displayed on the tabletop surface performed mostly one-handedly and was somewhat effortful. Moreover, prior studies revealed that users are still willing to use paper in combination with digital media on tabletops. In sum, prior studies independently investigated the usage and affordances of the way people interact with digital versus physical media.

In response to the aforementioned shortcomings of previous works, we conducted an exploratory user study in which we investigated the usage and interaction patterns when *concurrently* working with both types of documents on one surface.

2.1.2 Document Organization in Desktop Environments

There exists a large body of research investigating the management of digital documents in desktop computing systems. Although in this chapter we propose a novel user interface for hybrid document management on tabletops, we revisit related works that studied spatial layout and aggregation concepts for desktop computers. While not directly relevant, we believe that prior desktop systems that support transient grouping or piling of elements are good sources of inspirations for the design of the PaperTop interface.

Early studies have investigated and identified the advantages of piling versus filing in office desktop settings [Malone 1983, Whittaker 2001, Mander 1992]. The seminal work by Malone [Malone 1983] documents a series of interviews with ten knowledge workers on how they organize information on their desks and in their offices. He found that files and piles are the most salient ways people organize their workspaces. Files are units of elements that are explicitly titled and arranged, whereas in piles, the individual elements are loosely arranged and are not organized in any particular order. It was found that organizing information and filing them were cognitively difficult tasks. Based on findings, he suggested that digital systems can help knowledge workers by simplifying the filing of information and by providing intelligent aids for categorizing and retrieving information.

Later Whittaker and Hirschberg [Whittaker 2001] compared two ways of structuring information in office environments and found that piling is faster and easier to maintain than filing. They found that piles served as visual reminders and recent information is more accessible to people. Similarly, Mander et al. [Mander 1992] conducted a study to investigate how people deal with the flow of information in their workspaces. Their findings suggested piling and filing as the most common ways of organizing information. Contrary to filing, piling provided a casual lightweight way of organizing documents, which involved less overhead. Based on their findings, they proposed a design of a desktop interface for organizing documents in the form of piles. Users can then interact with piles by performing a set of gestures using the mouse pointer, such as gesturing sideways on a pile, which will spread out its element for the purpose of browsing.

The pile metaphor discussed above served as a foundation for many desktop interfaces to assist the management of digital documents. BumpTop [Agarawala 2006] is a good example of such systems, which aimed to bring more realism to today's rigid desktop computers by adding physical simulation and using piling instead of filing as the fundamental organization structure. In this way, virtual documents can be dragged and tossed around in a realistic fashion. Moreover, a set of pen-based interactions and visualization techniques enabled users to create and manipulate piles of virtual objects. Dynapad [Bauer 2005] is an extension to zoomable interfaces [Bederson 1996] that provided spatial tools for managing digital photos and an iconic representation of documents. It represents the objects in open piles on a zoomable Pad++ [Bederson 1996].

Bubble Clusters [Watanabe 2007] is an interface for manipulating spatial aggregation of digital objects. Based on spatial relationships, it automatically aggregates GUI objects into a group and visualizes a bubble around the objects (cf. figure 2.1c). To pick a hidden object, users can double click on the group to spread out aggregated objects. The bubble metaphor also exhibits hysteresis effects to avoid

(a) BumpTop [Agarawala 2006]

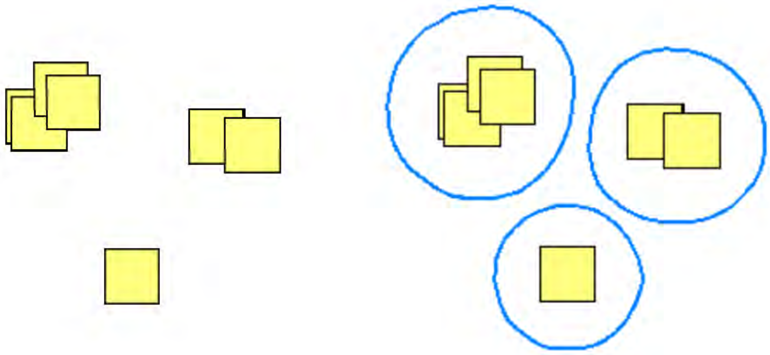

(b) Bubble Clusters [Watanabe 2007]

Figure 2.1: Desktop interfaces that organize information using spatial layout and aggregation

unwanted occurrences while merging or splitting groups. An evaluation with twelve users showed that the bubble cluster interface improved performance in a simple icon grouping task and an ink relocation task compared to standard folders and lasso selection, respectively. It was postulated that the bubble metaphor is suitable to represent transient structures and encourages further experiment in organizing documents.

Based on the findings of studies in this area of research, we can conclude that structuring documents in the form of piles is a promising way of organizing documents. This is because

- piles provide a spatial aggregation of a number of documents that represent a loose group structure
- piles' creation and manipulation require less cognitive overhead and their items are easily accessible to the users

- piles can also be spatially arranged in a space and convey different meanings to the user based on her location

Moreover, the results of the systems that support piling metaphor in their design showed that it improves task performance compared to the systematic filing metaphor. Furthermore, the bubble metaphor [Watanabe 2007] turned out to be a promising visualization for piles as their rotund shape indicates a casual and transient structure and compactly encloses objects they contain.

2.1.3 Hybrid Tabletop Systems

Interactive desks and tabletops enable users to manipulate physical and digital media on the same surface. This enables a novel interaction paradigm that closely integrates and transforms across bits and atoms. Therefore, it is important to review prior systems, focusing on supporting users on such *hybrid* physical-digital tabletop systems.

The seminal work of Wellner is the first tabletop system that combined the analog world of paper documents with the digital world of workstations through an augmented desk [Wellner 1993]. As illustrated in figure 2.2a, it consists of a real physical desk enhanced to provide some characteristics of a digital workstation (PC) such as projecting digital images down onto the desk or paper documents. Using an overhead camera, it is capable of capturing physical information and interactions of the user either with a stylus or the fingers. Correspondingly, a computer-driven projector is also mounted above the desk, allowing the system to project electronic objects onto either the desk surface or physical paper documents. In this way, users can for example do some physical number crunching by copying physical numbers from a real paper and pasting it on a digital calculator projected onto the desk [Wellner 1991].

Similar in nature and setup, several successor systems further improved the Wellner seminal work in terms of tracking fingers and hand gestures [Koike 2001, Robinson 1997], and proposed using an augmented desk for specific scenarios for example, in computer-supported learning [Mitsuhara 2010].

In the last decade, technological advances in multitouch displays have sparked a new surge of interest in supporting the co-habitation of physical and digital objects on the table. Recent tabletop computers are equipped with high-resolution cameras that enable real-time recognition and tracking of touch as well as physical objects using, for instance, fiducial markers. Based on the types and roles of the physical objects, research that took place in the last decade in this field can be categorized

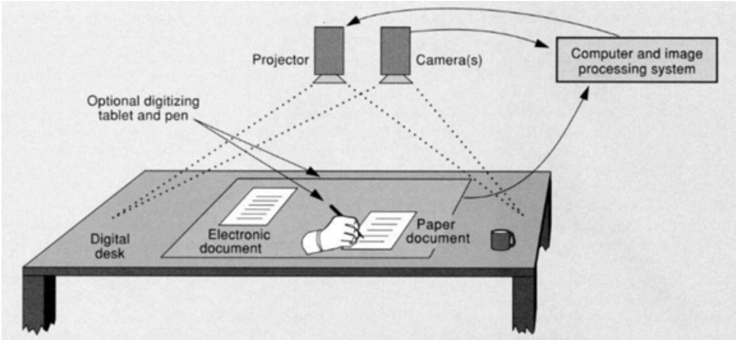

(a) Schematic view of DigitalDesk

(b) Digital calculator

Figure 2.2: The DigitalDesk by Pierre Wellner [Wellner 1993]

into two groups.

The first group considers systems that investigate the use of physical objects as tools to support manipulation of digital objects [Ullmer 1997, Jordà 2007, Kirk 2009, Terrenghi 2008, Hancock 2009]. A common characteristic of these interfaces is that they provide a tight coupling between physical and digital parts of the interface. Moreover, the physical components on their own have no function (utility).

ReacTable (cf. figure 2.3a) is a good example of such systems: the interface uses a set of physical objects as handles to mix, play, and create music on a tabletop computer. Users can also interact with the digital part of the interface represented on the tabletop surface (such as modifying a parameter around a physical handle) using multitouch gestures. In this vein of research, others also investigated issues such as design considerations of [Kirk 2009] and the effect of tangibility and physicality [Terrenghi 2008] in such hybrid interfaces. Their findings indicated that the tangible component of hybrid interfaces promotes bimanual interaction [Terrenghi 2008], is more precise [Hancock 2009], and provides eyes-free control and can reduce digital clutter [Kirk 2009]. These findings guided the design of the PaperTop interface.

The second stream of recent research encompasses prior studies that aimed at supporting true hybrid settings on tabletops. Typically, such settings consist of a set of everyday physical and digital objects like documents, a laptop, books etc. resembling the normal workspace of knowledge workers. A number of systems focused on transforming information between physical and digital domains [Rekimoto 1999, Haller 2006, Hartmann 2010].

Augmented Surfaces [Rekimoto 1999] is an augmented, spatially continuous workspace that consists of an interactive tabletop and wall allowing users to smoothly exchange

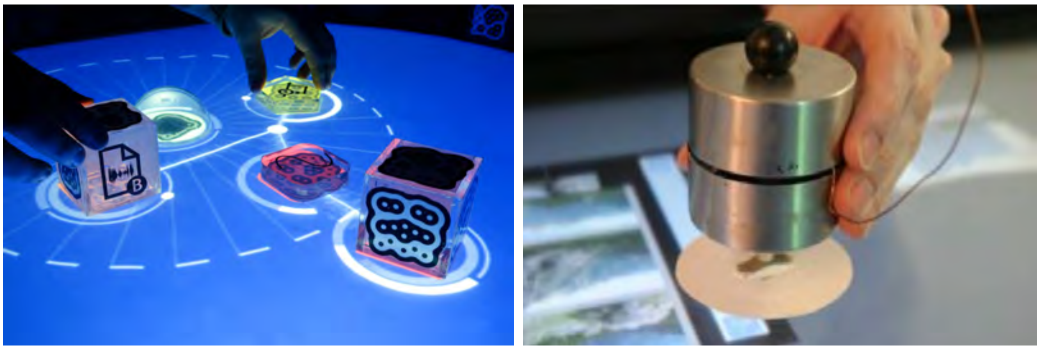

(a) Reactable [Jordà 2007] (b) Physical Handles [Terrenghi 2008]

Figure 2.3: Hybrid tabletop interfaces incorporating tangible objects as handles for manipulating digital content

and share information between their personal devices, tabletop, wall, and other physical objects. Knowing the physical layout of all objects placed in this workspace, users can move information in between personal computers and physical objects in a normal manner by dragging with a pointing device (e.g., the trackpad of a user's laptop). Similar in setup, Haller et al. designed an augmented reality tabletop system for supporting face-to-face collaboration among designers [Haller 2006]. The setup consisted of an interactive tabletop and wall on which users can interact using a stylus input. The system integrates users' virtual information as well as real objects, such as paper augmented with top-projected digital content. Digital content can be easily created on a paper document and seamlessly transferred from paper to the tabletop and even to other devices, such as a laptop.

Pictionaire [Hartmann 2010] is an interactive tabletop system to support design collaboration across physical and digital artifacts. It allows designers to seamlessly move information across physical and digital realms. The information on paper-based sketches are captured using a high-resolution overhead camera. The other way around (digital to physical) is realized through putting digital imagery on a physical tablet or notebook so that users can add annotations or create hand-traced copies of digital images. An evaluation with 16 professionals showed that Pictionaire promoted more collaboration compared to conventional analog and digital practices.

Considerable research has also explored creating and hyperlinking everyday physical and virtual objects on interactive surfaces [Everitt 2008, Klemmer 2001, Steimle 2009b].

DocuDesk [Everitt 2008] is an interactive tabletop that supports knowledge workers in retrieving the state of their tasks, which span the paper and digital

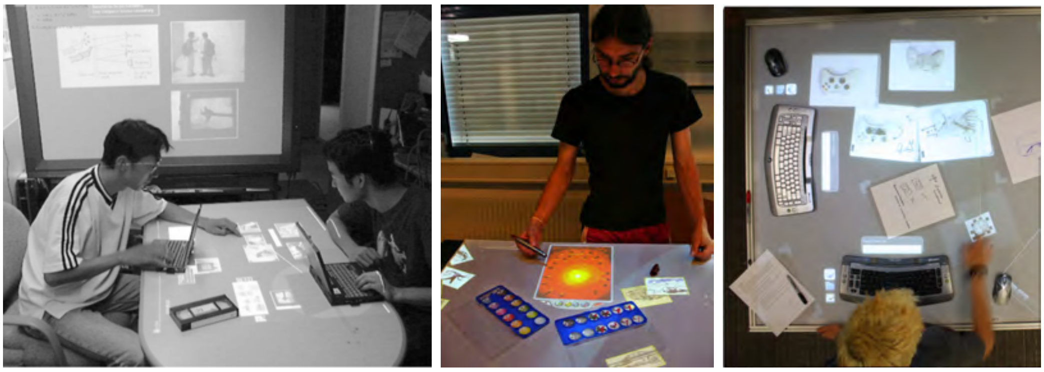

(a) Augmented Surfaces [Rekimoto 1999] (b) Shared Design Space [Haller 2006] (c) Pictionaire [Hartmann 2010]

Figure 2.4: Tabletop systems supporting real workspace involving everyday physical and virtual artifacts

worlds. It facilitates the creation of many-to-many relationships among paper and digital documents. The system recognizes paper documents through an overhead camera. Using a stylus, a user can draw a line to establish a link between a paper document and a digital one.

Likewise, CoScribe [Steimle 2009b] is a concept and a system that tightly integrates paper and digital documents on a tabletop surface assisting users in collaborative annotating, linking, and tagging of both digital and physical documents. It utilizes the same digital pen and the same interactions in both types of documents. Among the other techniques offered in CoScribe, users can create a hyperlink between physical and digital documents by marking the respective regions on both documents using the same digital pen. A first evaluation of CoScribe showed that it enhances both work performance and user satisfaction.

The Designers' Outpost [Klemmer 2001] is an interactive wall for collaborative website design that incorporates Post-it notes and digital inks. Users can add Post-it notes and stick to the wall. The system then recognizes and updates the knowledge on the board. Using a stylus, a user can create links among Post-it notes. The system keeps track of Post-it notes and their links so that users can easily re-position them on the board. The evaluation with 15 website designers underscored that the system feedback should not interrupt the designers' flow state and should only react upon explicit users' actions. Moreover, their findings revealed that there is substantial merit in a system that supports the coexistence of physical and virtual artifacts.

Based on the literature analysis presented above, we conclude that interactive tabletops support the natural, direct, and intuitive manipulation of both physical

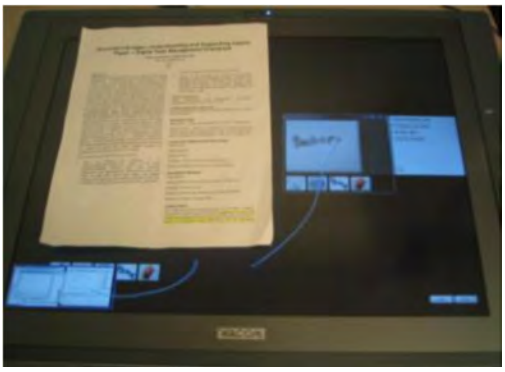

(a) DocuDesk [Everitt 2008]

(b) CoScribe [Steimle 2009b]

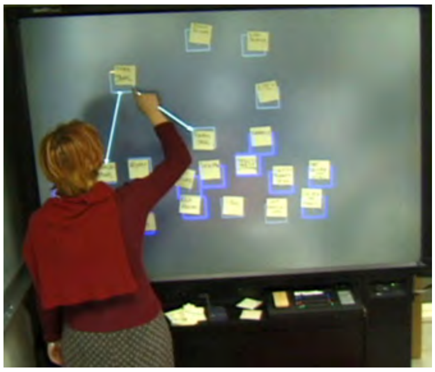

(c) The Designers' Outpost [Klemmer 2001]

Figure 2.5: Interactive surfaces supporting hyperlinking between physical and virtual documents

and digital objects on one surface. Therefore, it is promising to use them for supporting activities involving both physical and digital artifacts. Recalling findings of [Sellen 2003] that postulated *novel technology should work in concert with paper media*, we strongly believe that interactive tabletops can inherently support working with paper-based media and, at the same time, allow users to directly manipulate digital media with a greater spatial proximity.

2.1.4 Summary

We analyzed three different areas of the related work. In section 2.1.1, we reviewed prior work that studied the affordances of physical and digital media. We identified that paper offers several key affordances that are missed in the digital domain and that the design of digital systems should work in concert with paper-based media. Moreover, none of the prior studies in this field have focused on the combined use

⊕	Paper-based documents offer key affordances: flexible organization in space, bimanual interaction, strong haptic feedback.
⊕	Piling is a lightweight and efficient way of organizing documents.
⊕	Bubble visualization was found to be a promising visualization for transient grouping.
⊖	Previous studies examined affordances of physical and digital media independently.
⊖	Previous hybrid tabletop systems have not focused on supporting hybrid piling and grouping of documents.
⊖	Previous systems track physical objects bounded to the 2D surface of tabletops.

Table 2.1: Summary of the state-of-the-art analysis. ⊕s indicate important findings from the literature analysis that we considered in our interaction design, and ⊖s are lacking features that are addressed with our contributions in this chapter

of both types of media.

Section 2.1.2 presented an overview of desktop interfaces supporting digital document management. It was found that piling is a more lightweight way of organizing documents compared to filing. There also exist a number of desktop interfaces that support pile-based digital document management. Among others, the bubble cluster [Watanabe 2007] interface showed a promising visualization approach to support aggregation of objects due to its natural behavior and unobtrusive visualization.

We then review previous hybrid tabletop systems. It was shown that tabletops are suitable platforms that alleviate the manipulation of both types of media on such surfaces. However, there exist a few interfaces that support a hybrid physical-digital setting on tabletops. In particular, we found no interface supporting hybrid grouping and piling on tabletops. Moreover, most current hybrid tabletops track physical objects using a rear-mounted camera, which limits the use of objects to the 2D surface of the tabletop. Other tabletop systems, such as those of [Everitt 2008, Wellner 1993] that use an overhead camera to track paper sheets are also confined to the 2D space of tabletops.

Table 2.1 summarizes our findings from the examination of the related work, discussed above.

2.2 Understanding Hybrid Media Usage Patterns on Tabletops

As it was discussed in related work, although tabletop systems increasingly support activities involving both physical and digital media, there still exist a number of interaction design issues that need to be addressed. This section is a first step toward an effective integration of physical and digital media on tabletops, in which we aim to systematically assess spatial and behavioral interaction patterns when both types of media are grouped and piled on a tabletop work surface. Our main objective is to study the affordances and trade-offs involved in how users spatially arrange and group items. To this end, we conducted an exploratory study in which we emulate a hybrid tabletop setting. In general, with the aim of this study, we address the following research questions:

1. How do users concurrently manipulate physical and digital media on one surface?
2. Which salient behavioral patterns emerge from working in hybrid settings?
3. What is the impact of physical occlusion?

Note that results and findings related to the third question are presented in the next chapter where we address problems of physical occlusion.

2.2.1 Method

Participants

We recruited ten volunteer participants for the study five of whom were female. In terms of handedness, eight were right-handed and two were left-handed. Their ages ranged from 24 to 46 years, with an average of 32. All participant were experienced knowledge workers, from both technical and nontechnical backgrounds such as computer science (7), psychology (2), jurisprudence (1). All but two had little or no particular experience with interactive tabletops, but all were familiar with multitouch input, for instance, on smart phones. No compensation was provided.

Study Setup

The study was conducted in our lab environment. We used a 130 x 105 cm custom-built interactive tabletop of with a display size of 100 x 60 cm. This is representative of the space available on a typical desk. The rear projection had a full HD resolution of 1920 x 1080 pixels. Rear projection is currently the most common form of tabletop display. The tabletop is back-illuminated with diffused infrared light, and a PointGrey Flea2 CCD Camera is mounted inside the tabletop to observe the screen. The camera has a maximum resolution of 640 x 480 pixels at 80 frames per second, which provides a responsive and real-time touch and fiducial marker recognition.

The participants could interact with digital documents using typical multitouch gestures for moving, rotating, and enlarging or shrinking individual documents. Printed documents could be placed and manipulated on the display surface and also on the surrounding non-display areas of the table. The setting is shown in figure 2.6.

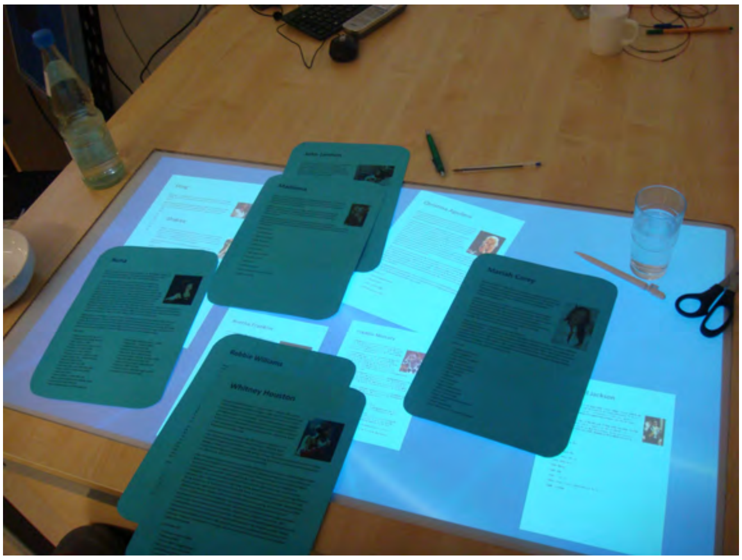

Figure 2.6: Study setup

Materials

Participants were given 12 single-page documents. Each document contained textual biographical information about a popular musician. We used textual documents instead of photos (as used in other studies) for two reasons. First, working with documents is a more common and thus more representative knowledge work task. Second, since documents are typically larger than photos, the possibility of occluding digital items increases. To create a hybrid media setting, six of the documents were printed on paper (A4 size), and the remaining six were displayed on the tabletop display. In a typical work setting, users might not have a common set of information items arrayed across digital and physical media but rather different types of items in each (e.g., paper articles or notes used while consulting web information).

We moreover used a few nonmedia objects, such as a water bottle, a few pens, a scissor, and a bowl to generate more realistic hybrid setting as well as to study the impact of occlusion generated from arbitrary everyday objects. In each task, physical documents and digital documents were initially presented as two separate adjacent piles positioned at the center of the tabletop. All participants used the same set of documents.

Procedure

Each user participated in a single-user session lasting about one hour. The specific tasks we asked participants to perform are common in knowledge work with printed and digital documents: a grouping task (inspired by [Terrenghi 2007]), and a search and comparison task. These tasks enabled observation of component activities, such as selecting, reading, comparing, and moving items.

Each session started with introducing conventional gestures for interaction with digital documents (move, zoom, rotate) to participants. They then performed the grouping task and were asked to browse the documents and create three groups: those describing artists they like, those they don't like, and those they are unsure about. The search and comparison task required finding all albums that were released in the same year as one specific album by one artist. Therefore, participants had to compare all documents with one specific paper document and note down the names of the albums. To help compensate for learning effects, different criteria for grouping were used. While performing the tasks, subjects were asked to think aloud. After each task, semi-structured interviews were conducted. All sessions were videotaped, and two observers took field notes.

2.2.2 Results

We coded 6.75 hours of video recordings for relevant behavior, iteratively refining the coding scheme. This resulted into a set of categories that describes behavioral patterns. We are particularly interested in the affordances of the combined use of physical and digital media. This includes the general patterns of use as well as hybrid grouping behaviors, which are described below in turn.

2.2.2.1 Spatial Patterns of Use

We are especially interested in how tabletop space is used in a hybrid setting. An initial analysis of the video recordings showed that participants distinguished between territories for working with (e.g., reading, comparing) and for storing items. As also reported by [Scott 2004], the boundaries between these areas were quite flexible. We coded usage locations for each participant and aggregated this data. Figure 2.7 shows this aggregated data for the hybrid search and comparison task. Darker locations depict areas that were used by more participants. The wooden frame of the table is indicated by a gray border that encloses the white display area. The dark gray circle indicates the position of the user.

Figure 2.7 (a-c) shows the areas used for working with items. Note that physical and digital items tend to be used at different tabletop locations (graphically summarized in figure 2.7 [a-b]). Physical items nearer to participants were used more often and more commonly to their left. At times they were not directly on top of the display but on the edge of the table or even partially jutting off the table's surface. In contrast, digital items were more tightly clumped, primarily positioned to the right and farther away from the user than the physical items. Note also that the working area for physical items involved a larger portion of the table, while the area for digital items was almost exclusively confined to the right of the table display. A reason might be that the haptic feedback of physical documents affords cognitive off-loading, whereas the direct-touch manipulation of digital media requires visual attention. This fact could explain why the participants delegated macrometric tasks to the non-dominant hand, and micrometric ones to the dominant hand [Guiard 1987].

Figure 2.7 (d-f) depicts the areas used for storing items (i.e. locations of temporarily placed items for later use and of final placements). While the working areas were centered more in the middle of the table, storage is situated toward the outer left and right edges. Also, unlike working areas, storage area placement patterns are similar across physical, digital, and hybrid uses. Note the clear difference in the use of the tabletop area between storage and working, especially for physical and digital items. Again placement of physical items extends beyond the table's surface. It is also important to notice that in contrast to working areas, physical and digital storage areas were co-located, with the placement of physical and digital items overlapping.

Figure 2.7 g and h shows not only that physical items were used on the surface of the table but that there was also significant manipulation of these paper items above the table. We classified usage as being above-the-table if participants held a physical item (or collection of physical items) above the table for a period of time longer than was required to pick it up and directly place it elsewhere. The illustration further shows that items in the air were held nearer the user and more to the left than the physical items on the table's surface. We observed that when participants wanted to interact with a physical document, in the vast majority of cases they first selected a pile (45 cases) or an individual item (256 cases) by picking it up and holding it in their hands above the table. Only then would they engage in further interactions, such as reading or comparing information. Much less frequently, participants did not pick up the item and interacted directly on the surface with the pile (8 cases) or the individual item (20 cases).

We also observed the different behaviors of participants while they engaged in both types of documents in the search and comparison task. Participants in this task had to compare information contained on a physical item with information from all other items. While all participants interacted with digital items by enlarging them, strategies with physical documents varied. Three participants permanently held the paper item above the surface with the non-dominant hand. Seven participants placed the physical item on the table's surface near the digital items. This contrasts with how participants compared only physical items, where the majority held items in the air above the surface.

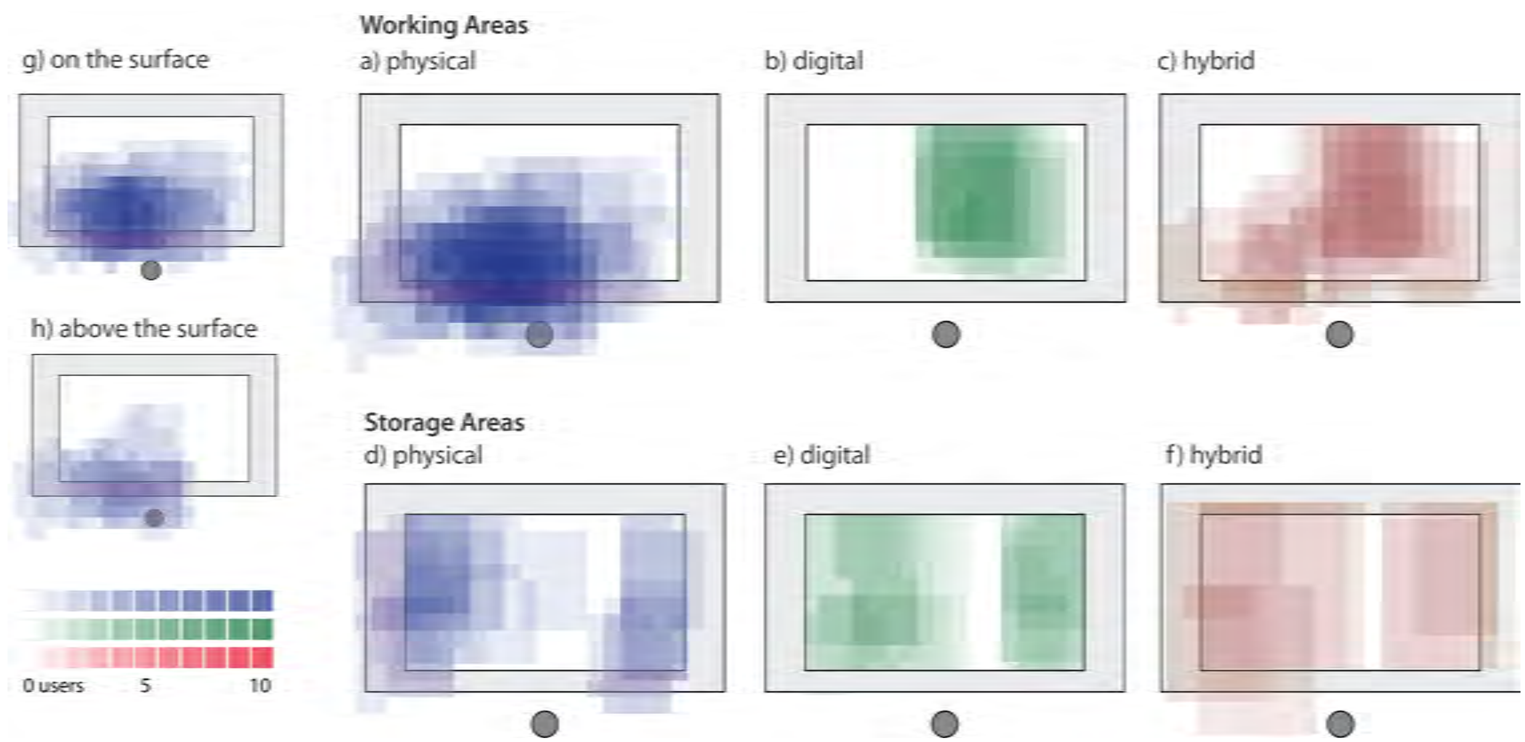

Figure 2.7: Activity maps depict areas where items were used on the hybrid surface in the hybrid search and comparison task. Darker areas signify use by more participants [Steimle 2010b].

2.2.2.2 Hybrid Grouping

We analyzed how participants interacted with groups that contain both physical and digital items. We were particularly interested in knowing whether users prefer representing a hybrid group as two separate, possibly adjacent groups, each containing only physical or only digital media (spatially-separated representation), or if they prefer arrangements that integrate both types of items at the same place (spatially integrated representation). All participants grouped physical and digital items in a single spatially integrated representation, a layered arrangement that we call a *hybrid pile*. This is an arrangement of largely or entirely overlapping physical and digital documents. Figure 2.8 shows examples of hybrid piles created by participants in the grouping task.

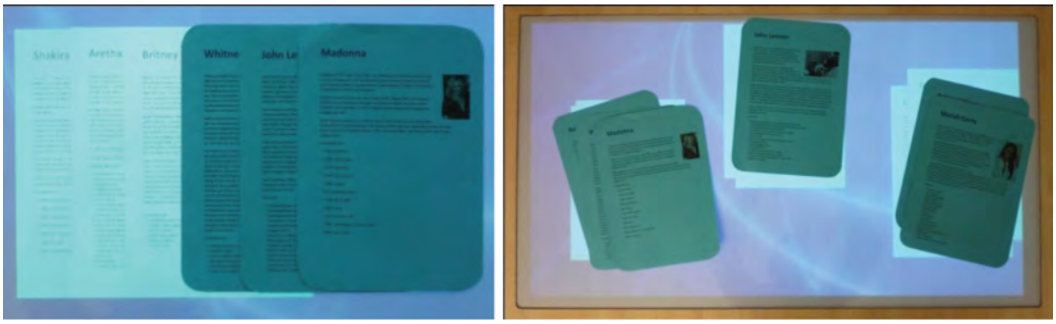

Figure 2.8: Examples of hybrid piles on the tabletop. Left: Spread-out representation, which affords getting an overview and comparing pages. Right: Result of grouping task created by one participant showing three examples of hybrid piles [Steimle 2010b].

The spatially integrated hybrid representation enabled participants to express a close relationship between physical and digital documents. Comments during interviews also underscored that participants preferred this spatially integrated hybrid representation to other representations. For example, P9 explicitly mentioned that "having all [digital and physical] items in a pile makes sense." P3 stated, "I don't want digital documents to be [automatically] slid out" when placing physical documents on them. Due to the high degree of spatial proximity of overlapping physical and digital media, the overall gestalt [Koffka 1935] of a hybrid pile clearly expresses that the elements of the group belong to one conceptual group. Several participants mentioned that once a hybrid pile is formed, the system should provide feedback about its state.

The creation and sequential browsing of hybrid piles appeared fluid and dynamic. All but one participant created hybrid piles by first placing digital items on top of each other and then placing physical items on top of the digital items. To browse hybrid piles, all participants sequentially picked up the topmost item of the pile or dragged it away, starting with the physical before going to the digital items.

Hybrid piles are more static than physical piles. Pure physical piles were frequently moved, whereas we observed no instance of an entire hybrid pile being moved. In addition, we observed frequent transformations of physical piles. Different pile organizations simplified specific practices. Placing all items neatly on top of each other saves screen real estate and supports interacting with the group as a whole (e.g., for moving and storing). In contrast, spreading out the items affords looking at the contents of several documents in parallel (e.g., for getting an overview and comparing items). In the pure physical condition, we frequently observed participants performing these transitions with one quick and intuitive bimanual movement. In contrast, in the hybrid case, we observed only two instances of such transitions. An obvious reason might be the fact that each digital item of the pile had to be displaced individually to transform the arrangement.

2.2.3 Summary

In this study, we have explored the affordances and trade-offs involved in how users spatially arrange and group documents on the tabletop surface. The results revealed a spatially integrated use of physical and digital documents. In particular, participants frequently stored groups of documents in a hybrid pile arrangement. Participants were willing to physically occlude digital contents in order to better manage their workspaces and make semantically meaningful collections (e.g., hybrid piles). On the other hand, when it was needed, interaction with digital occluded documents in hybrid piles turned out to be demanding and ineffective. Table 2.2 summarizes the main results of this study, from which we infer several implications presented in the next section.

2.2.4 Requirements

Based on the findings of the study, we infer four essential requirements as rationale for an interaction design to support hybrid media usage on tabletop systems:

R1. Support for both physical and digital information grouping

Our study demonstrates once again that people can work very flexibly with collec-

General Usage Patterns
– Physical documents were mostly used above the tabletop surface.
– Physical items were placed nearer the user than digital items and were less clumped.
– Interaction with piles of physical documents was found to be quick, easy, and intuitive.
– Haptic feedback of physical documents affords bimanual manipulation and cognitive offloading.
– Digital interaction was performed mainly by the dominant hand and required visual attention.
Hybrid Grouping
– *Hybrid piles*: the most frequent form of organizing both types of documents.
– Unobtrusive and continuous feedback about the state of a hybrid pile was missing.
– A fast and easy way to change hybrid pile representations was missing.
– Tidy piles afforded a compact form factor for ease of moving the entire pile.
– A spread-out (fanned-out) representation was used for getting a quick overview.
– A juxtaposed representation allowed for detailed comparing of physical and digital items.

Table 2.2: Summary of the results from the hybrid media usage study

tions of physical documents. Users easily interact with and move groups of physical items to accomplish tasks and manage workspace organization. Hybrid piles without better technological support lack the ease and flexibility of moving and rearranging that traditional piles of paper documents afford. Therefore, it is particularly important that tabletop systems provide easy natural mechanisms for users to create groups of hybrid items and enable interaction with them as single entities. More precisely, easy manipulation of hybrid piles on the tabletop surface – such as the creation, re-organization, adding, or removing of items – should be supported. Furthermore, the state of the pile should be continuously transmitted to the users. This means that it should be visible what documents belong to which pile. Recalling the prior hybrid tabletop systems discussed in 2.1.3, none of the prior systems have supported hybrid information grouping on tabletops.

R2. Support for fluid transitions between pile representations

We frequently observed that users transformed piles of physical items into another arrangement to assist the task being performed. For example, a “tidy” pile affords interaction with the group as a whole, whereas a juxtaposed or partially overlapping arrangement affords an overview of the items as well as reading or comparing them. With physical items, transformations were fluid and typically performed with one bimanual movement. Similarly, fluid transitions between different representations should be supported for digital and hybrid groups. The digital members of a hybrid pile could, for example, automatically imitate arrangements that result from moving physical members. As we have seen in the related work (2.1.2), several desktop interfaces supported such pile transitions using a set of explicit mouse gestures. However, to the best of our knowledge, there existed no tabletop interface that addressed this issue.

R3. Support for physical document use above and around the surface
Physical items not only were used on top of the display surface but were also frequently placed on the margins of the table (even jutting off the surface), as well as frequently picked up and held in the air above the table. The physical interaction space is larger than the digital interaction space, extending in all three dimensions. As a consequence, hybrid tabletop systems that plan to support the use of physical items should track them not only on and directly above the surface but also on and above the table edges and even in front of the table. The review of hybrid tabletop systems in 2.1.3 showed that most of the tabletop systems track physical items with fiducial markers using a camera below the table. This confines the tracking of items to the 2D surface of the tabletop. With our study, we showed that tracking should cover areas beyond the tabletop – for example, our findings suggest areas in front of the display and to the side of the table associated with the user’s non-dominant hand.

In table 2.3, we summarized the requirements discussed above. We also showed which of them are met by previous work from section 2.1, then listed the corresponding contributions of this thesis to overcome limitations of previous works.

These requirements are addressed in the present chapter, in which we describe the PaperTop interface and visualization concepts that support hybrid grouping on tabletops. While explaining PaperTop concepts and techniques, we refer back to requirements that are supported.

	Requirements	Supported by previous work?	Contribution of PaperTop
R1	Support for both physical and digital information grouping	○	PaperTop interface allows for integration of both media types on tabletops.
R1a	Allow for flexible and intuitive interactions with a hybrid pile	○	PaperTop leverages tangibility of physical documents and do not need explicit commands for hybrid pile manipulation.
R1b	Provide for lightweight and continuous feedback about the state of a pile	◐	PaperTop provides rich visual feedback using the bubble metaphor.
R2	Support for fluid transitions between pile representations	◐	PaperTop recognizes physical pile transitions and automatically rearranges digital documents.
R3	Support for physical document use above and around the surface	◐	PaperTops identifies and track physical documents on, around, and above the surface.
R3a	Reliable tracking of physical document in 3D space above the surface	○	PaperTop can robustly track paper documents, even those that are partially occluded.

Table 2.3: Overview of our requirements and the extent to which they are covered in the previous work. ○and ◐show whether the requirements have partially been addressed in the previous work or not, respectively.

2.3 PaperTop: An Interface for Hybrid Media Piling on Tabletops

In this section, we propose *PaperTop* – a novel user interface concept and set of interaction techniques to address challenges of hybrid physical-digital piling on interactive tabletops. In the following, we first present the underlying interface concepts, then a set of concrete interaction techniques for hybrid piling on tabletops, in which we exploit physical interaction with several paper documents for manipulation of digital ones.

2.3.1 Underlying Interaction Concepts

The results of the study have shown that hybrid piles are an intuitive and adequate representation for (unordered) groups of paper and digital documents on tabletops. However, current hybrid tabletops lack the intuitiveness, directness, and flexibility that we enjoy when interacting with groups of paper documents. As we have seen in 2.1, prior hybrid tabletops incorporate either specific custom-made physical items as controls for manipulation of digital items (e.g., [Jordà 2007]) or everyday objects for better supporting tasks that span across physical and digital realms (e.g., [Hartmann 2010]). In the former case, physical and digital items are tightly coupled in the interface, and physical items do not have any function on their own. On the contrary, physical everyday objects – such as printed documents – are *first-class objects* that users can exclusively work with. Inspired by both types of physical functionalities, we propose exploiting physical documents for interaction as well as controlling digital items to support the interaction with hybrid piles. More precisely, we utilize the physical pile as a tangible control for the entire hybrid pile, thus enhancing hybrid piles with the ease of interacting with paper. Compared to the traditional tangible UIs [Ishii 1997] and hybrid surfaces [Kirk 2009], this has the following advantages:

- The use of paper document as an everyday physical object that users frequently work with eases interaction with the user interface as a whole. This makes the interface more intuitive and minimizes the user's effort to learn how the user interface functions (R1, R1a)

- A paper document has a dual function in our design: it acts both as a tangible control for the digital documents and as a first-class object, since it consists of documents that have value on their own (cf. figure 2.9, left).

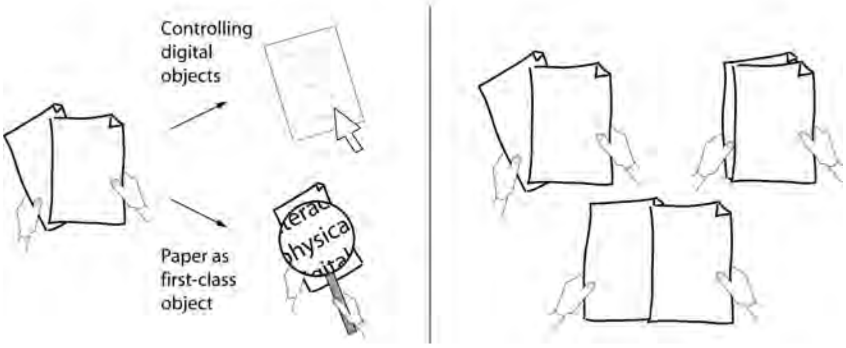

Figure 2.9: Paper pile as a tangible control to support manipulation of hybrid piles [Steimle 2010a]

- Typical tangible user interfaces offer tight one-to-one coupling between their physical and digital components. In our design, the physical control does not consist of one single object, but of several sheets of paper contained in a physical pile (cf. figure 2.9, right). In this way, we can leverage various arrangements of several sheets of paper as tangible controls to manipulate digital objects (R2). This offers more degrees of freedom for controlling the digital part than one single object.

Although using paper piles as tangible controls open up new possibilities and maintains the intuitiveness and flexibility of real-world interactions (R1a), it poses challenges for the interaction design. In particular, the dual function of the physical objects makes it difficult to determine if users interact with them in order to manipulate the digital part (tangible control), or they are interested in the physical documents per se and does not want to alter the digital part (first-class object). Another challenge is that it is not clear when a document belongs to its group, when it is disjointed, and, foremost, when a hybrid group is formed. Of course one primitive solution is to use explicit commands to mitigate these challenges. One basic solution is to use typical rubber band or lasso selection techniques that we use in today's desktop interfaces and extend them to select a number of objects in a group on multitouch tabletops. However, these solutions need explicit commanding to group objects before performing any operation with the entire group (e.g., relocation or split).

On the contrary, the interface design in PaperTop is particularly inspired by the natural behavior of people managing temporary or transient information in physical world. We found in prior studies [Malone 1983, Mander 1992] as well as in our own study (section 2.2) that people use a spatial layout to represent the semantic relationship among objects (documents). Objects placed closer to each other are semantically related, and farther-placed object are unrelated. Therefore, in the PaperTop interface, documents (whether physical or digital) that are placed close to each other are automatically recognized as a group and PaperTop visualizes the results as a bubble surrounding the objects [Watanabe 2007]. In this way, grouping does not require any explicit commands (R1b). As we discussed in the related work section (2.1), the bubble visualization offers several advantages [Watanabe 2007]:

- Its round shape represents a transient structure which encourages further experiments in organizing documents.
- Their shape is spatially efficient compared to a box or a convex hull, resulting in saving space on the tabletop surface.
- It resembles the behavior of a soap bubble, so users can easily understand and pick up its interaction style (R1a and R1b).

The bubble visualization provides clear feedback about which documents belong to which group (R1b). Moreover, the flexible form of the bubble provides for flexible rearrangements of the documents contained within the group. Two bubbles join and become one large bubble if they touch; in the reverse, one bubble can also be divided into two bubbles. Leveraging this flexible nature, we use bubbles to continuously provide unambiguous feedback about the hybrid grouping of documents and to allow easily and flexibly creating and rearranging groups (R1a and R1b).

In contrast to pure digital bubble clusters designed for interacting with a mouse [Watanabe 2007], our approach addresses hybrid interaction with both physical and

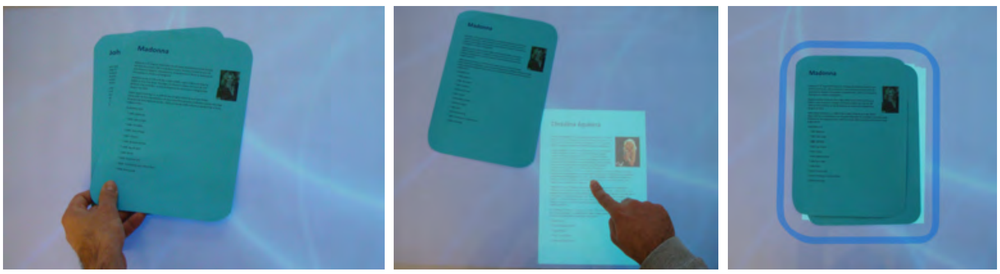

(a) (b) (c)

Figure 2.10: Pile creation by (a) placing physical on digital documents or (b) dragging digital under physical one. (c) A hybrid pile is formed [Khalilbeigi 2010].

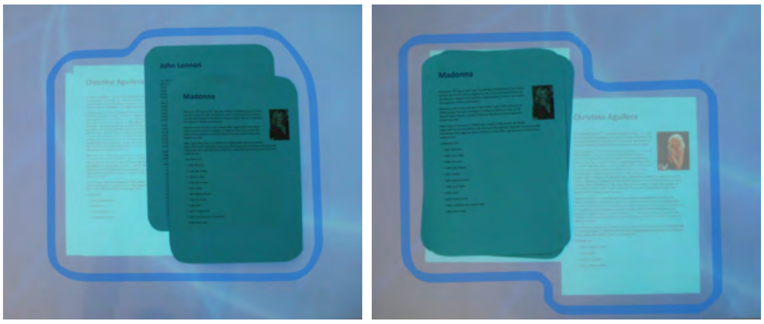

Figure 2.11: Flexible reorganization of the elements of a hybrid pile: the bubble adapts itself to the arrangement of the pile [Khalilbeigi 2010].

digital items. This allows for more flexible and implicit ways of interaction, such as using physical items as a graspable handle for the hybrid group. The bubble metaphor facilitates the main piling activities in hybrid settings, presented as a set of concrete interaction techniques in the following.

2.3.2 Interaction Techniques

Wrapping

By placing (or dragging) documents onto each other, either a hybrid group is formed or a new item is added to an existing hybrid group (cf. figure 2.10). A bubble is automatically displayed and visually wraps all documents contained in the same group (R1 and R1a).

Stretching

The bubble visualization supports very flexible arrangements of the documents contained in a group. Individual documents within a group can be displaced. Figure 2.11 illustrates two examples in which documents of piles are placed one beside the other for reading and comparison purposes while maintaining the hybrid group structure. The bubble flexibly adapts its shape to the new contours of the group (R1b). In contrast to the work by Watanabe et al. [Watanabe 2007], it is not only possible to interact with one single item at a time, leveraging bimanual and tangible interactions, but the user can also grasp and move several items very easily.

Tidying Up

In our study, we observed two main representational forms of groups:

- a tidy pile of documents, which affords interaction with the group as a whole

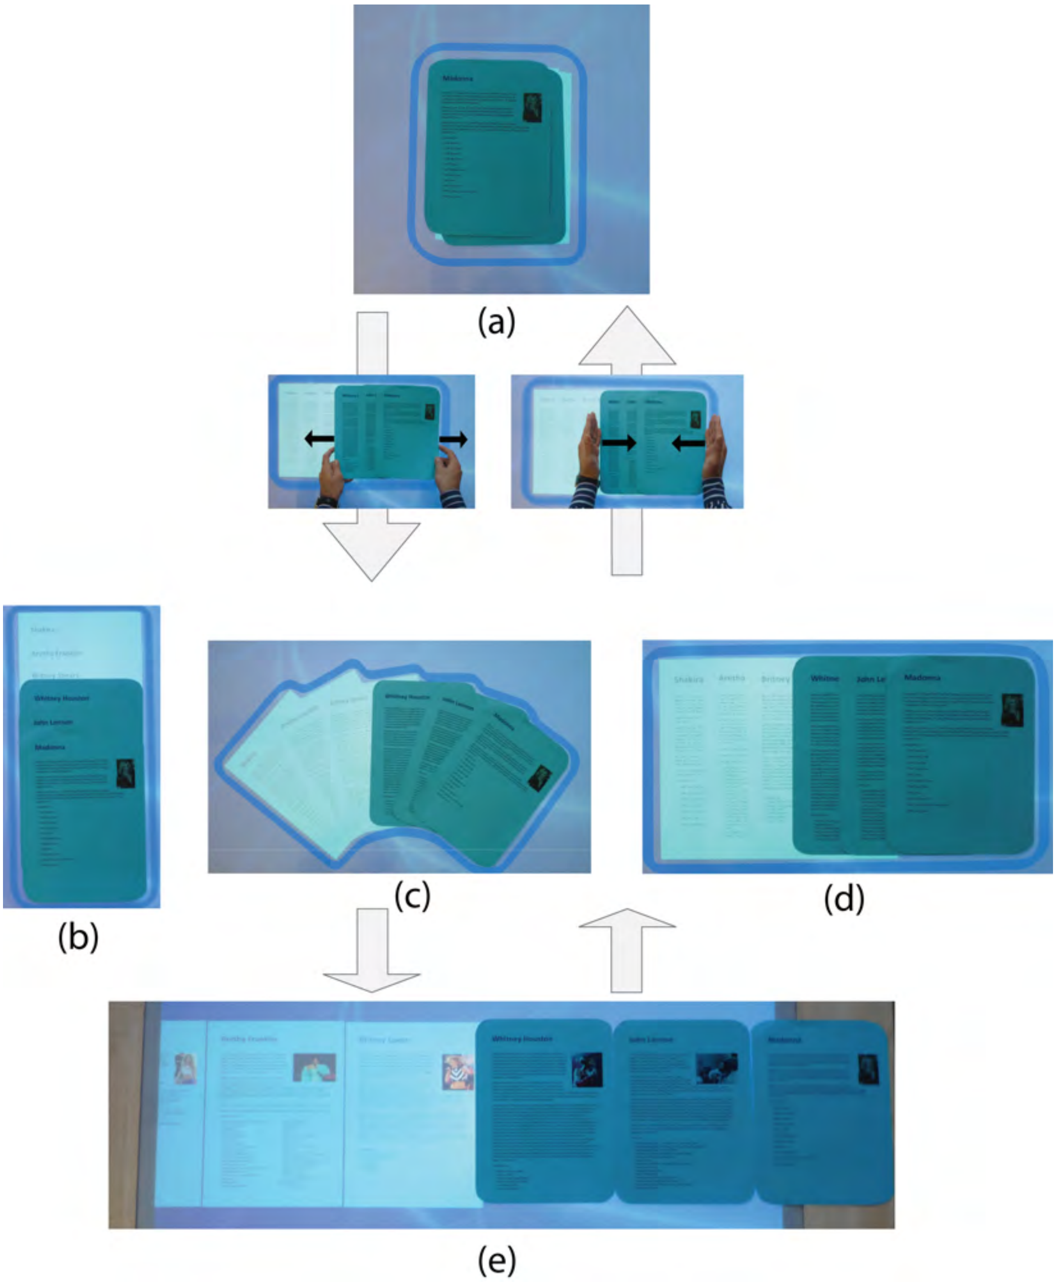

Figure 2.12: Fluid transitions for hybrid piles, from (a) the compact form to intermediate forms – (b) top to down, (c) fan out, and (d) left to right – to (e) a full juxtaposition [Khalilbeigi 2010]

(e.g., moving)

- a juxtaposition or partially overlapping arrangement, which affords getting an overview of the documents, reading, and comparing them.

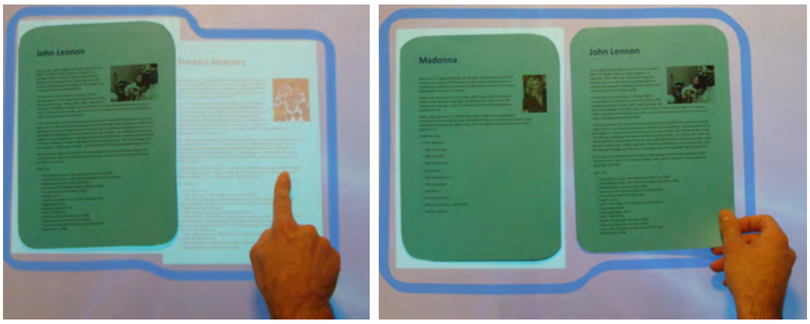

Figure 2.13: Moving individual physical or digital documents without affecting the entire pile [Khalilbeigi 2010].

Groups of paper documents allow fluid transitions between both representations. Our design leverages the many degrees of freedom of the tangible control to offer fluid transitions between hybrid group representations. In order to arrange all documents of a hybrid group in a particular representation, the user has to arrange only the physical part in this manner. The digital documents of the group are automatically relocated and rotated to form the same representation. Hence, these hybrid transitions are as easy to perform as with pure paper (R2). They are depicted in figure 2.12. These transitions can be used for getting an overview of all documents of a hybrid pile. Moreover, they allow quickly switching between different forms of interaction (e.g., moving a group vs. comparing the documents).

(a) (b) (c)

Figure 2.14: Transferring interaction to relocate an entire hybrid pile by using paper documents as tangible controls [Khalilbeigi 2010].

Transferring

The ability to move and drag documents on the tabletop plays an important role in optimally managing the workspace. In our study, we observed that physical piles

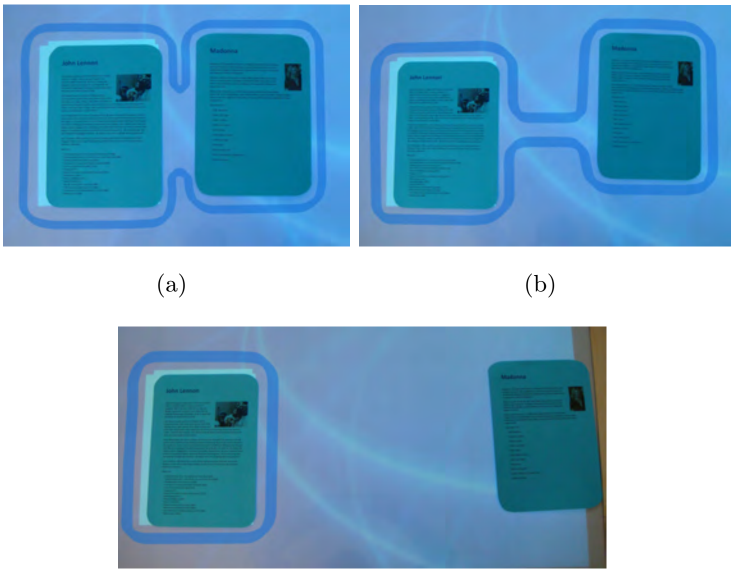

(a) (b)

(c)

Figure 2.15: Removing a document from a hybrid pile. The soap bubble continuously visualizes the state of the pile [Khalilbeigi 2010].

were relocated very frequently by moving or picking-and-dropping. In contrast, digital and hybrid groups were never moved because suitable functionality was missing. We want to leverage paper as a tangible control for easy moving of hybrid groups. However, this concept raises a conceptual challenge. If the user moves documents, it is not clear if she wants to interact with these documents, using them as first-class objects (e.g., for reorganizing the pile as described above), or if she wants to use them as a tangible handle for the entire pile. Our design resolves this ambiguity as follows:

- If individual documents, but not the entire physical pile, are moved, the documents are treated as first-class objects (i.e., they are rearranged without affecting the remainder of the pile). This interaction is depicted in figure 2.13.

- If all documents of one entire physical pile are moved or picked and dropped (cf. figure 2.14) together, they are considered as a tangible handle of the entire hybrid pile. Automatically the digital part of the hybrid pile is relocated according to the movement of the physical part. In the special case that the user wants to move all documents of the physical part at the same time, but without affecting the digital part, she can hold one of the parts at a fixed position while moving the other part away. Then both parts can be moved separately. They can be easily reconnected by simply placing them onto each other as of creation.

Deletion

If a document of a group is relocated to be far from the other documents of the group, the single bubble of the group gradually splits into two separate bubbles. As depicted in figure 2.15, while dragging the document away, the user gets continuous feedback to indicate when the document gets separated from the group, because the connecting part of the bubble will grow smaller. If the distance exceeds a threshold, the document(s) is/are removed from the group and form(s) a separate document or group.

Table 2.4 summarizes the five interaction concepts supporting interaction with hybrid piles.

Name	Purpose	Description
Wrapping	Creation of hybrid piles	Bubble visually wraps pile items by putting physical item(s) on digital ones or moving digital item(s) under physical ones.
Stretching	Flexible reorganization	Bubble continuously adapts its shape while stretching the pile items.
Tidying Up	Fluid transitions between pile arrangements	Digital items are automatically re-arranged when users change the arrangement of physical ones.
Transferring	Moving hybrid piles	The digital part moves with when users reposition the physical part of a hybrid pile.
Deletion	Removing an item	Moving away individual items of a hybrid pile so that the bubble gets disconnected.

Table 2.4: Summary of interaction concepts included in the PaperTop interface

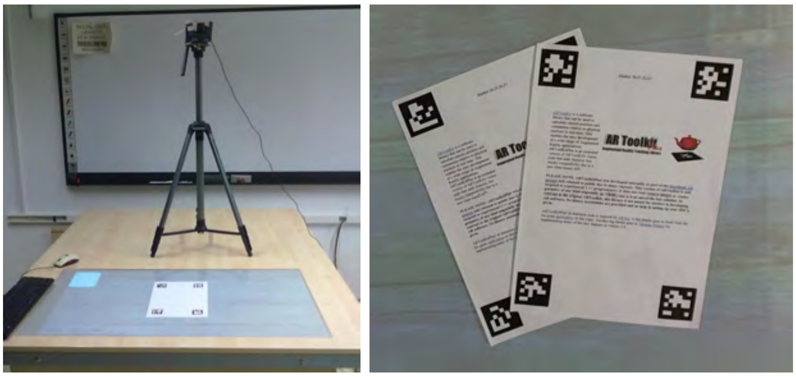

(a) (b)

Figure 2.16: (a) Hardware setup of PaperTop, (b) multiple sheets of paper, each augmented with four fiducial markers for reliable tracking

2.4 Implementation

We have implemented a PaperTop prototype and its techniques running on the tabletop system introduced in 2.2. Here we first describe the hardware setup and then the implemented algorithms for tracking paper sheets in the 3D space above the tabletop (R3) as well as the bubble visualization technique.

2.4.1 Hardware

Figure 2.16a shows the hardware setup of the PaperTop prototype. The main component is the tabletop (as introduced in the study setup 2.2), which runs all software components. It was found in our study that physical documents are mainly manipulated above or out of the tabletop screen area (either at the user's hand or placed partially or fully on the tabletop rims) and therefore it is not possible to recognize and track paper sheets using the camera inside the tabletop. In the PaperTop prototype, a top-mounted, high-resolution webcam is used to identify and track multiple sheets of paper in the space above the tabletop. We used fiducial markers and printed four markers on each corner of the paper sheets in order to robustly track them (R3a). In this way, if a corner of a paper sheet is bent or covered with the user's hand or another paper sheet (cf. figure 2.16b), the tracking algorithm can still extract the paper location and orientation using other markers printed on the paper. Other document recognition and tracking approaches such as those based on visual features [Liao 2010] do not perform very reliably. This is because of the fact that in our setting, it is very likely that paper sheets partially occlude each other, resulting in ambiguous shapes or visual appearances.

2.4.2 Algorithms

Software of PaperTop consists of two main components: i) document tracking and clustering, and ii) visualization. For reasons of efficiency, the document and clustering component runs on a separate PC, and the visualization component runs on the tabletop computer. All components communicate with each other using MundoCore [Aitenbichler 2007] in real time.

2.4.2.1 Document Tracking and Clustering

In order to support the interaction techniques proposed in 2.3, the system needs to recognize (ID) multiple sheets of paper and track them in the 3D space above the tabletop surface. This means that for each paper sheet, we need information on the position (x, y, z) as well as the orientation $(yaw, pitch, roll)$. This is realized in a three-step pipeline described below.

Step1: Marker Tracking
We used ARToolkitPlus [Wagner 2007] as the marker tracker, which is freely available under the GPL open-source license. It extracts the six degrees of freedom (6DoF) pose information of markers from the 2D planar camera image in real time. As reported in Wagner's work [Wagner 2007], ARToolkitPlus is capable of tracking ten markers in 8.3 ms, running on a system with an Intel 3.0 GHz Core Duo CPU. Our system prototype uses a PC with an Intel Core i3 4.0 GHz CPU and 4GB RAM, which provides sufficient computational power to support real-time interaction techniques. In this step, raw camera images streamed from the webcam are fed into the tracker, and it returns the ID, position, and orientation of each detected marker in the scene. Based on the returned ID of markers, the system then classifies which markers belong to which paper sheet.

Step2: Document Tracking
Based on the individual marker information, in this step the position and orientation of individual paper sheets are calculated. We first calculate the position of a paper sheet based on xy values of the respective markers attached to it. The z value of the marker tracker is not precise enough to determine how paper sheets are piled. Therefore, we only calculate the Z index of paper sheets which means how they are stacked on each other. It is calculated based on the visibility of the four markers of

each paper sheet to the camera. For example, the Z indices of the underlying and overlying paper sheets in figure 2.16b will be 0 and 1, accordingly.

Since the paper sheet can be partially bent while users hold it in their hands, orientation values of the marker on that side will be largely different from the other markers. Therefore, instead of averaging orientation values, which may result in an incorrect 3D pose of the paper, we calculate the median value of the markers' orientation. In this way, the 6DoF information of each paper document is determined.

For digital documents, a tracker component keeps track of the 2D position of digital document displays on the tabletop. Through a calibration process, we calculate a transformation matrix so that the positional information of both types of documents is transferred to one unique coordinate system. This is needed for hybrid piling and visualization, presented next.

Step3: Pile Creation and Tracking

In this step, all pile-related logic is calculated. This means that based on the proximity of documents (both physical and digital) and the visibility of the four markers of each paper sheet, the system reasons about all hybrid pile interaction techniques – such as creation, reorganization, moving, removing, etc. In our implementation, each document conceptually forms a pile in which only one item exists. If two documents (piles) overlap by more than 80%, they merge into one pile. In order to support flexible reorganization of piles, the items remain attached to a pile if its documents overlap by less than 80% and are not moved far away from each other (the distance between the documents' center points should not exceed 25 cm). Once a document exceeds the threshold, it will be removed from the pile. If the entire pile is moved (i.e., four markers of the top-most document in the pile remain fully visible to the camera), the digital documents will follow the physical ones.

To support fluid transitions between various hybrid pile representations (cf. figure 2.12), the system needs to first recognize various representational forms of paper sheets and then, automatically rearrange the digital part. To do so, it examines positional information (x, y) of the physical documents' center points as well as the direction in which physical documents are arranged. This is done by analyzing the z index of documents in the pile. In this way, if a set of physical documents has approximately the same x- or y- coordinate values, then vertical or horizontal representations are recognized and the digital documents will be arranged accordingly. For the fan-out representation in addition to the (x, y) coordinate of center points and direction, a radius is calculated describing the circular form of the fan-out arrangement (cf. figure 2.12c).

2.4.2.2 Visualization

The visualization component running on the tabletop computer takes a set of documents belonging to a group and generates the bubble visualization that surrounds the documents. In order to do so, we have implemented a 2D version of the metaballs and isosurfaces concept of blobby shapes proposed by Blinn [Blinn 1982]. Each metaball function is defined as

$$f(x_1, y_1) = \frac{Radius^2}{(x - x_1)^2 + (y - y_1)^2} \tag{2.1}$$

in which x and y are the center points of the metaball and x_1 as well as y_1 are coordinates of arbitrary points in the space. This equation is based on the equation for calculating the strength of an electrical field in science, which provides the largest value in the center of the metaball and will drop off quickly as the distance from the metaball gets larger and larger. In our system. we defined each document (be it digital or physical) using the metaball model in which the center point and the maximum distance from the document edge to the center point determine the center point and the radius of the associated metaball function accordingly. For multiple documents, the potential fields of adjacent documents are summed to generate a smooth curved boundary. Since the metaball function defines an isosurface in the 2D planar space, we then need to define the maximum and minimum thresholds that extract an isocontour, which surrounds the documents. The isocontour is defined as follows:

$$\text{min. threshold} \leq \sum_{k=0}^{n} Metaball_k(x, y) \leq \text{max. threshold} \tag{2.2}$$

In the current implementation, we empirically determined the optimal threshold values by performing several experiments with a number of various digital and physical documents. Other approaches – such as the standard marching-squares method [Lorensen 1987] – can be envisioned for more effective isocontour detection.

2.5 Evaluation: Initial Expert Feedback

We evaluated our interaction concepts for hybrid piling using the prototypically implemented system in an early user-feedback session. The session was conducted with HCI researchers recruited from the local university department. The main objective of the study was to get a first impression if the techniques are conceptually sound and to see how the experts as knowledge worker interact with hybrid piles using our techniques.

2.5.1 Method

We recruited four HCI experts to evaluate the techniques in individual sessions. The study took place in our lab environment, which is an open space containing ordinary desks and one digital tabletop, as described in 2.2. The participants were all male, 28 years of age, and have participated in our first study. They have an average of four years' professional experience in the field of HCI, particularly interaction design.

The study setup consisted of three sheets of paper; each was tracked with four markers printed on each sheet corner. This resulted in a total of 12 markers to be tracked with our system, which in some pretesting found out to be a reasonable number to support interactions in real time. We used three digital documents manipulable using typical multitouch gestures. Although all participant participated in the first study, we first introduced the tabletop system and the context of the study. Then they were asked first to familiarize themselves with the tabletop and find a comfortable sitting position. We then introduced the interaction techniques for hybrid piling. Then they were given enough time to use the system and try out the techniques. We asked participants to think aloud while playing around with the system. As a data gathering method, we observed participants and at the end of each session conducted a semi-structured interview. We transcribed the data and analyzed salient quotes. Each session lasted about 30 minutes.

2.5.2 Results and Discussion

All four participants appreciated and easily understood the hybrid piling techniques. They all were able to quickly pick the interaction style of bubble visualization. P2 said, "After few interactions with hybrid piles one can easily understand what its [bubble] job is." Compared to the first study setting where no hybrid support was provided, all participants felt that the both types of documents act more in concert: "both documents are somehow aware of each other and work more harmonized." From participants' comments we found that since our system provide a lightweight feedback and a set of coherent interactions based on natural behavior of working with paper, it integrated well with established practices of working with paper. Thus participants all were pleased while working with the system.

Despite the fact that one of our main goal was to design implicit interactions, three participants wished to have explicit interactions for accomplishing some tasks. Three participants stressed that spatial proximity used to define an aggregation or disaggregation of documents in group is not always a meaningful criteria. For instance, P4 said that in some cases you need an explicit split command for disjointing a document from the pile *in place*. Another participant mentioned that browsing of pile can be achieved by performing some sort of explicit gesturing on the bubble visualization. A concrete example is tapping on it and dragging out to see an overview of digital documents. In general, we observed that participants found the bubble visualization somewhat *static* and preferred more interactions with digital documents through gesturing with the bubble.

Fluid transitions between pile arrangements was well received by all participants. In particular, they liked the horizontal arrangement either partially or fully juxtaposed. They found it of great value: "It often happens that I arrange my paper documents in horizontal way on my normal desk and again push them together to make a neat pile" as P1 stated. The fanned-out arrangement was found to be quite impractical. Nevertheless, we found that the high degree of tangibility provided by physical documents allows participants to easily change the arrangement of hybrid piles.

2.6 Conclusion

Interactive tabletops are becoming more popular and widespread due to the new interaction paradigm that they offer. Among others, enabling representation of digital media in a life-size way and providing a direct manipulation through multitouch gestures are the most prominent advantages. As a result, researchers and practitioners started exploiting them for information-based activities and providing digital support for knowledge workers. Because of their table like form factor that allows placement of physical next to digital media they become particularly of interest to bridge the gap between these two realms.

In this chapter we approached bridging this gap by exploring the *concurrent* use of physical and digital media on interactive tabletops. Through a user study we looked at how people naturally work in such hybrid settings and what are the emerging challenges. We observed that people organized both types of documents as they do in the real world: spatially grouped documents in form of piles that are semantically related. Hybrid piles were the most salient form of organizing document in hybrid tabletop setting. People frequently rearrange them in order to browse or obtain an overview of their documents. Moreover, it was found that interaction with physical documents occurred mostly on the space above the surface. Despite the easy and lightweight interactions with physical documents, interaction with hybrid piles turned out to be cumbersome.

To mitigate these challenges, we contributed PaperTop, an interface consisted of set of interaction and visualization concepts for facilitating common knowledge work practices such as piling or grouping in hybrid tabletop settings. In our design we particularly leveraged the tangibility offered by physical documents to provide intuitive means for controlling digital items. Based on spatial proximity of documents, the techniques enabled easy creation, moving, deletion and rearranging hybrid piles. We also designed an unobtrusive visualization technique that support the casual and transient nature of piling. We leveraged the metaphor of bubble to monitor the state of hybrid piles to the users. Once a hybrid pile is formed, a bubble like visualization will surround the group items. It flexibly adapts its shape when people engage with the pile.

To evaluate the effectiveness of the PaperTop interface we implemented all concepts in a prototype system. It supports 3D usage of physical documents and keeps track of all documents. Using this prototype we could evaluate our concept in early feedback study in which we invited several HCI experts to examine the concepts. The results of the study was promising and all experts appreciated the use of paper documents as means for controlling the digital part of hybrid piles. It was found to be intuitive and easily understandable.

Chapter 3

Managing Physical Occlusion on Tabletops

Contents

In the previous chapter, we addressed challenges emerging from the integration of physical documents with interactive surfaces and proposed a set of physical interaction concepts to support basic knowledge work practices such as grouping or piling. While physical documents are the most essential and prominent element of knowledge work, they cohabit with many other *everyday objects* on tables. Tables in the real world frequently clutter with people's personal and work-related objects – like a coffee cup, a wallet, a paper tray, etc.

The present chapter explores the integration of *physical everyday objects* with interactive surfaces. As mentioned in the introduction of this thesis, this is an essential step toward making horizontal interactive surfaces, *real tables* that can be seamlessly integrated into everyday life. However, the presence of physical everyday objects on today's tabletop computers mars display and interaction with digital contents. Physical objects may hide or partially occlude digital items, create issues because users are unaware of the presence of occluded items, make access difficult, and restrict screen areas for display and interaction.

In this chapter, we focus on physical occlusion problems. Based on an exploratory study, we contribute to the understanding of how people deal with occlusion of screen contents by physical objects and identify emerging challenges. Our analysis reveals a dual character of occlusion involving both inconvenient and desirable aspects. People are willing to physically occlude digital content in order to better manage their workspace, switch the task context and make meaningful collections such as the hybrid piles addressed in the previous chapter. When needed however, awareness about and access to the occluded objects turns out to be cumbersome. Moreover, physical objects hinder interactions with digital objects and thus, organizing digital workspace becomes a tedious task.

To address these challenges, we propose *ObjecTop*, an occlusion management framework mitigating the physical occlusion problem on hybrid interactive tabletops. The framework consists of a set of interface concepts supporting awareness of, access to and organization of digital objects while appreciating the positive semantic aspects of occlusion. These concepts are informed by the arrangement of physical objects on the tabletop surface and enable users to access occluded digital items. The occlusion-aware interface concepts are iteratively designed and evaluated in two *design-implement-analyze* cycles.

In summary, the main contributions of this chapter are

- exploring the impact of physical occlusion in hybrid tabletop settings on user's performance
- establishing a set of design requirements for hybrid tabletop systems mitigating physical occlusion problems
- design and implementation of an occlusion management framework providing

access-, awareness-, and organization-supporting interaction concepts

- examining the occlusion-aware interaction concepts through two user studies iteratively

The structure of this chapter is as follows. In 3.1, we review previous work that investigated the problems of occlusion in 2D windows desktop environments. This is followed by discussing prior studies aiming to cope with occlusion caused by people's hands or forearms on touch- and pen-enabled displays as well as by physical objects in hybrid tabletop environments. In 3.2, we present our occlusion-related results from the exploratory study conducted in the previous chapter. In addition, we conducted a second study in which we examine use of everyday objects on interactive tabletops.

Based on the findings of both studies, a set of design requirements are listed in 3.2.4 that set foundations for our occlusion management framework presented in section 3.3. We then report user studies in which we evaluated our concepts and techniques in 3.5. We close the chapter by a conclusion (3.6) and outlining future directions.

Contribution Statement: Most of the work presented here is based on and has been published in [Khalilbeigi 2013, Khalilbeigi 2012b]. I am the first author on these publications. I have initiated and lead the project. My co-authors have also contributed significantly. Master students, Patrik Schmittat and Jan Riemann, have built and implemented many aspects of the ObjecTop system. My supervisors, Jürgen Steimle, Max Mühlhäuser, and James D. Hollan, have contributed to the design of the system and on writing the papers.

3.1 Related Work

In this section, we review the state of the art in three research areas pertaining to the occlusion problem in different settings. First, we revisit previous systems that dealt with the problem of 2D digital occlusion in desktop computers. This is followed by discussing studies that investigated occlusion caused by the hands and fingers while using stylus-based or direct touch input modalities. Finally, the most related to our work is those that coped with the problems of physical occlusion on interactive surfaces. In the following we discuss previous work in each of these research areas in turn.

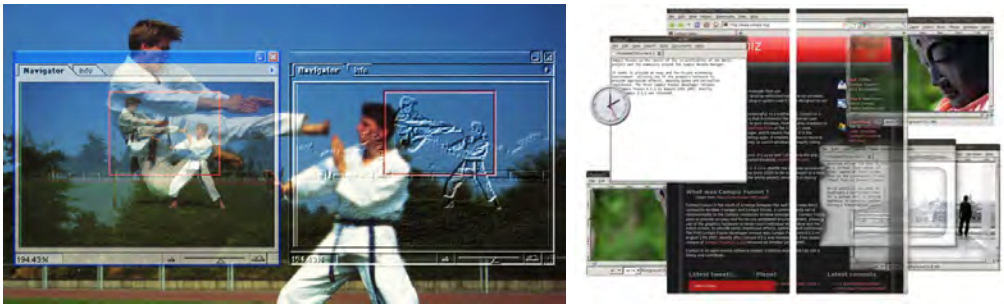

(a) (b)

Figure 3.1: Transparetizing approaches to reveal hidden 2D content: (a) alpha blending and multiblending techniques used on the right- and left-side windows accordingly [Baudisch 2004], (b) content-based transparentizing technique [Waldner 2011].

3.1.1 Occlusion of Digital Objects

Digital occlusion has been a topic of research for a long time. There is a large body of research investigating occlusion problems in 2D windowing and drawing desktop environments. While not directly related to our focus i.e. hybrid tabletop environments, we describe how a similar problem in mainstream 2D windowing systems is mastered.

In a 2D window-based system, such as Microsoft Windows, the desktop workspace of users can easily get cluttered with an array of application windows stacked on top of each other. In order to access and interact with the underlying windows, users either need to scale down the windows or retrieve them through special widgets like the Windows taskbar or the MaxOS Dock. Several window management approaches avoid this problem by cascading or tiling the windows in empty space on the desktop. Other approaches avoid digital occlusion by placing windows in regions with the least occlusion [Bell 2001].

Another stream of research to cope with digital occlusion has focused on altering the level of transparency of windows in order to reveal the underlying content. Alpha blending [Porter 1984] makes the occluding digital content semi-transparent to provide hints about the underlying contents. Harrison et al. [Harrison 1995] examined a semi-transparent tool palette to allow two windows to be displayed on the same piece of screen space and thus, reduce the need for switching between overlapping windows. They found that 50% transparent palettes can greatly improve workspace visibility without degrading icon selection performance.

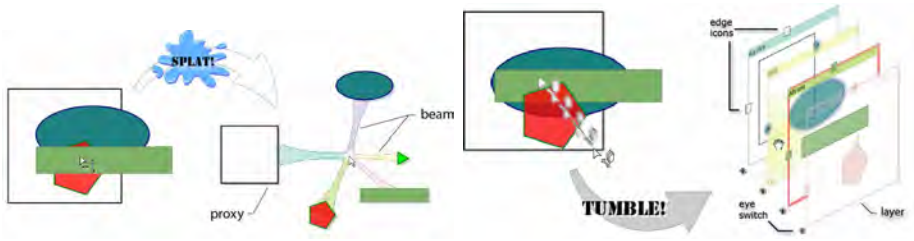

(a) Splat (b) Tumble

Figure 3.2: Splat and Tumble techniques presented by Ramos et al. [Ramos 2006]

Similarly Baudisch et al. [Baudisch 2004] proposed a multiblending technique that uses a vector of blending weights, one for each class of features rather than a single constant transparency value used in the alpha blending technique. Both techniques are illustrated in figure 3.1a. Other approaches took the content of Windows into consideration and designed more effective techniques that apply transparency to, for instance only unimportant (white) window regions [Ishak 2004] or unimportant regions that are visually less salient [Waldner 2011] which is shown in figure 3.1b.

Researchers also investigated more sophisticated approaches. Ramos et al. [Ramos 2006] presented two techniques named Tumbler and Splatter to support accessing and manipulation of occluded content in 2D drawings. The Splatter technique (cf. figure 3.2a) reveals the occluded content by spreading out the stacked objects while preserving their shape and size. On the other hand, Tumbler (cf. figure 3.2b) is basically a layered visualization which mainly preserves the z-value of stacked 2D objects on each other. Authors have evaluated techniques in controlled experiment and found that both techniques can co-exist and complement each other based on user's task and scenario.

While these approaches provided practical solutions for desktop computers to be used with keyboard and mouse, they can rarely be employed for hybrid tabletop computers. This is mainly because of the fact that digital objects cannot be displayed on physical ones. Nevertheless they guided the design of our techniques. In particular, techniques presented in Ramos et al. [Ramos 2006] in which the underlying objects are uncovered by displacing from their original position to a nearest non-occluded location. Moreover, the use of proxies in the Splat technique (cf. figure 3.2a) for selecting and manipulating occluded contents highly inspired our proxy visualization described later on.

3.1.2 Occlusion by Hands and Fingers

The advent of touch and pen-based interaction modalities has naturally resulted in situations in which a portion of display content may get occluded by user's finger, hand and forearms. This gave rise to the same problems to that of physical occlusion on hybrid tabletops. Therefore it is important to analyze previous work in this very field of research and to set out design guidelines for our concepts.

Shift [Vogel 2007] addressed the problem of finger occlusion on small handheld touch screen devices by displaying a callout showing a copy of the occluded screen area and place it in a non-occluded screen area. The callout also shows a pointer representing the selection point of the finger so that users can guide the pointer into the target by moving their finger on the screen surface. The evaluation showed that users can select small targets with much lower error rates that on an unaided touch screen.

In an experiment Vogel et al. [Vogel 2009] showed that the hand and forearm can occlude large portions of the tablet display while using a stylus. Based on their analysis they presented a geometric model of the hand occlusion which can enable new types of occlusion-aware interaction techniques on tablets. Building upon this, Vogel et al. [Vogel 2010] proposed occlusion-aware interfaces which know what area of the display is currently occluded using a real time reconfigurable geometric model of the hand occlusion area on the display surface. Utilizing this knowledge, the interface can represent the occluded content in non-occluded areas using a bubble-like callout. In an experiment, the authors tested the technique in a simultaneous monitoring task and found that the occlusion-aware technique can successfully mitigate the effects of occlusion. The techniques include occlusion aware pop-ups, dragging, and hidden widgets depicted in figure 3.3.

Occlusion from the hand on interactive tabletops has also drawn the attention of researchers. Brandl et al. [Brandl 2009] present an adaptive menu placement method based on the handedness of users. Their technique recognizes the occluded area under the user's hand based on direct pen and hand tracking on the tabletop. A direction vector from the hand to the pen determines the correct orientation of the menu. In this fashion, they designed a circular menu that avoids occlusions caused by the user's hand.

Wigdor et al. [Wigdor 2006] proposed a two-sided interactive tabletop that receives touch input from both the top and bottom surfaces of the table. Using the bottom side of an interactive tabletop eliminates finger and hand occlusion while touching the surface. Vogel et al. [Vogel 2012] presented results of user studies

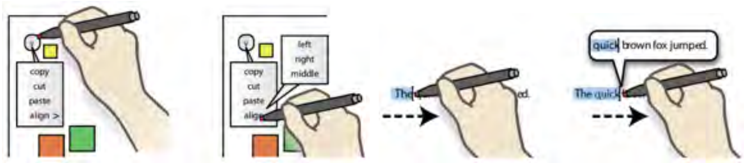

(a) Occlusion-aware pop-ups (b) Occlusion-aware dragging

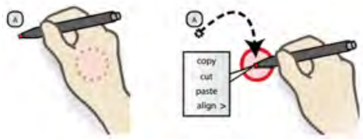

(c) Hidden widget

Figure 3.3: Occlusion-aware technique designs proposed by Vogel et al. [Vogel 2010]

in which they examine the shape of the hand and forearm occlusion on the multi-touch table for different touch contact types and tasks (tapping, dragging, rotating, and scaling). By analyzing the shapes of mean occlusion silhouette visualizations they found common characteristics across people, postures, and tasks. Based on their study results they propose occlusion-awareness templates for designers that use dimension-calibrated images to show areas that may be occluded relative to the expected contact centroid. These templates guide occlusion-related decisions – for instance, in placement of application widgets in non-occluded areas. Recently, Adachi et al. [Adachi 2013a] proposed using the forearm as an interactive surface that is very easy to access and an occlusion-free area for displaying context menu while working on tabletops.

In summary, we have seen that the finger, hand and forearm considerably occlude display content while interacting with bare fingers or a stylus [Vogel 2009]. We are also inspired by the general strategy to cope with physical body occlusion, which first identifies the occluded area and then visualizes a callout that shows a copy of the occluded screen placed in a non-occluded area. It was also found that despite distinct shapes, hand postures have common elements – like a hand blob and a protruding forearm, which mainly occlude display areas while doing dragging and tapping tasks [Vogel 2012]. These findings guided the design of our occlusion-aware techniques.

3.1.3 Occlusion of Physical Objects

Finally, the most related sets of previous work are those that investigated problems of physical occlusion on tabletop displays. In the recent decade, work is emerging that

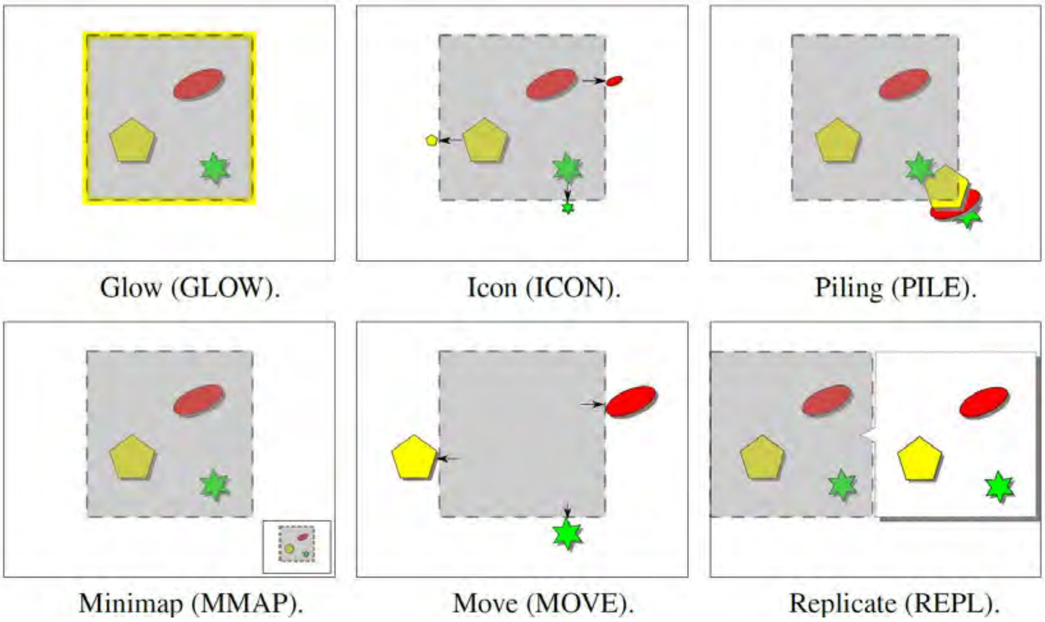

Figure 3.4: Overview of occlusion management techniques presented by Javed et al. [Javed 2011]. The large gray rectangle represents the footprint of a physical object placed on a tabletop display.

partially addresses the problem space of physical occlusion. Javed et al. [Javed 2011] presented six different techniques to address occlusion in a physical-virtual setting. The techniques are depicted in figure 3.4. Each of these techniques helps users deal with a single issue – such as identifying or accessing occluded digital objects. The techniques are evaluated through two experiments in a simulated hybrid setting using virtual occluders, but no physical objects. Experiments focus on guided and unguided visual search tasks with various numbers and sizes of occluded objects. Their results indicate that techniques that consume less space and provide less visual clutter are in general more promising. The study shed light on the characteristics of each individual technique and inspired the design of some of our techniques.

The goal of the SnapRail [Furumi 2012] interface was to alleviate the visibility and manipulability problems of virtual objects occluded by physical objects on tabletops. The interface is basically a rail widget that appears around the physical occluder. Once occlusion happens, occluded virtual items are automatically rearranged and snapped around the rail widget. Users can then browse the virtual elements along the rail widget. A preliminary user study of the SnapRail interface showed that it was easy to use and received positive feedback from participants.

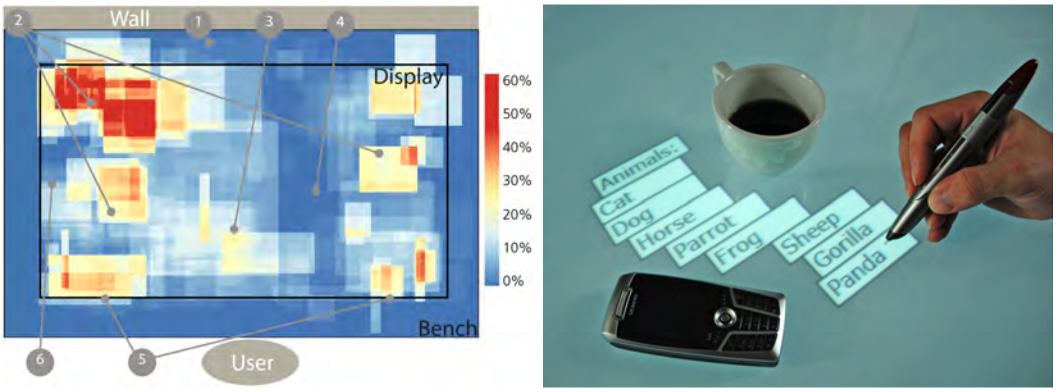

(a) (b)

Figure 3.5: (a) Five different zones found in Tabard's work [Tabard 2013] on which participants placed physical objects and (b) a user-drawn menu designed by Leithinger et al. [Leithinger 2007] using a stylus

Tabard et al. [Tabard 2013] investigated how users managed physical and digital objects during the 16-week deployment of a tabletop in a biology lab setting. The study was conducted on the eLabBench [Tabard 2012] tabletop system to support synthetic biologists. For the study, four participants used the eLabBench, covering 15 hours of bench work, extending over five sessions each lasting from 8 to 30 minutes. Based on analyzing video codes, their results revealed that on average, 15% of the tabletop screen was occluded at any point in time. They found five different tabletop zones where participants placed objects as illustrated in figure 3.5a:

1. long term storage: outside of tabletop screen, in which no occlusion occurred
2. storage: areas located at a far distance to users
3. observation: areas for short and transient observation
4. digital: where digital objects are mainly used
5. interaction: where small physical objects are frequently placed and used.

Based on their results, they suggested that occlusion management techniques should either limit distractions as new objects are put on the table or avoid occlusion rather than focus on reacting to it.

Other researchers have proposed approaches to occlusion avoidance for hybrid tabletop settings. Cotting and Gross [Cotting 2006] introduced environment-aware display bubbles that aim at significantly enhancing the flexibility, interactivity, and adaptivity of displays. They described a top-projection display metaphor that automatically projects free-form bubbles on a tabletop. Using a warping approach the display content is represented by multiple bubble forms that can elastically be manipulated.

Leithinger et al. [Leithinger 2007] investigated digital menu representation on mixed digital-physical tabletops. In contrast to automatic placement, they presented a user-driven approach in which users, using a stylus, draw menus of custom shapes around areas that may be occluded by physical objects (cf. figure 3.5b). In this way, users can simply draw an occlusion-free menu with less confusion. Their design also requires less technical computations for automatic recognition of physical objects. They proposed four different metaphors derived from natural ways of spreading information. They compared these techniques with traditional pie and pop-up menus in normal and cluttered environments. Their techniques showed a significant reduction of time in cluttered settings.

Freeman et al. [Freeman 2013] presented a technique to find uncovered areas on a tabletop surface to represent occluded content. Their technique first detects the footprint of physical objects atop the tabletop using a blob detection algorithm, and constructs a binary matrix that shows occluded areas of the display. Using this matrix, it is then possible to search for the largest unoccluded rectangular area that is suitable for presenting information.

3.1.4 Summary

Table 3.1 summarizes the examination of the related work discussed above. In this section, we analyzed related work in three different areas. We first revisited previous work that dealt with occlusion in 2D digital environments. Although not directly related to the challenges created by physical occlusion, these studies motivated us to use a proxy-like representation of occluded items as a means to aid ease of access. We then review prior studies that addressed challenges raised by physical body occlusion on touch- and stylus-based displays. It was found that an ample amount of display gets occluded while interacting with such displays. We moreover learned that the callout visualization is a promising way to visualize occluded content. Finally, we discussed previous works addressing problems of physical occlusion. We learned that techniques that consume less space and introduce less visual clutter are more promising for general tabletop applications. Our analysis of the literature in this field also revealed that occlusion management techniques should limit distractions

$\oplus$ A proxy-like representation of occluded items eases accessing of occluded objects.
$\oplus$ Callout visualization is suitable for displaying occluded content on occlusion-free areas.
$\oplus$ Techniques that consume less space and introduce less visual clutter are promising for general tabletop applications.
$\oplus$ Occlusion-aware systems should limit distractions caused by the system.
$\ominus$ In addition to supporting awareness of and access to occluded objects, an occlusion management system should support interaction with and organization of digital objects.
$\ominus$ Previous studies have not been examined in true hybrid settings with everyday objects.
$\ominus$ None of the previous studies compared their system with the natural way of coping with occlusion (i.e., moving or lifting physical occluders).
$\ominus$ Prior occlusion avoidance systems do not allow for intentional occlusion.

Table 3.1: Summary of the state-of-the-art analysis. $\oplus$s indicate important findings from the literature analysis that we considered in our interaction design, and $\ominus$s are lacking features that are addressed with our contributions in this chapter

caused by the system.

The present chapter significantly extends the previous work in three distinct ways. First, our contributions present an integrated interactive solution and techniques for hybrid tabletop settings that seamlessly supports *awareness*, *access*, and interaction with virtual objects as well as their *organization*. Second, we study our interaction concepts in a *true hybrid setting* using everyday physical objects. This enables assessing the influence of physical properties of various physical objects on the strategies users employ to cope with occlusion. Third, our evaluation compares the techniques with the natural way of coping with occlusion of everyday physical objects: *moving or lifting the occluder*. To our knowledge, this is the first such study. Furthermore, in contrast to the occlusion-avoidance approaches, we allow for intentionally occluding digital objects and at the same time provide occlusion-aware support.

3.2 Understanding Physical Occlusion on Tabletops

In this section, our goal is to address the impact of physical occlusion in hybrid tabletop settings and identify emerging challenges. Toward this goal, two steps are taken. First, we report occlusion-related results of the study described in the previous chapter (see 2.2). We analyze how physical occlusion influenced common activities and interaction with digital objects – such as selecting, zooming, or moving. Since the study involved only printed documents, we carried out another exploratory study incorporating various physical everyday objects, as our second step. We then compile findings of both studies into a set of design requirements that set the basis for our physical interaction concepts, presented later in this chapter.

3.2.1 Results of the Hybrid Media Study

As we mentioned in 2.2, one of our main objectives was to analyze the impact of occlusion on interaction with digital objects. As was expected, we observed that the presence of physical documents on the interactive surface restricted the interaction with digital objects. In the following, we present our results on how physical occlusion influenced the common activities of accessing (selecting), zooming, moving and grouping digital items. Note that descriptions of the study methodology, design, and data analysis are described in 2.2.1.

Accessing Occluded Documents

Occlusion was most salient in our data when users wanted to access a digital item partly or entirely occluded by physical items. In order to analyze strategies used for accessing digital items, we characterized each selection activity from the video recordings. We noted the strategy used, characteristics of the occlusion (percentage of occlusion, number of occluding items), as well as features of the surrounding tabletop surface (percentage covered by physical and digital items).

Figure 3.6 depicts the four strategies we identified and their frequencies. The most frequent strategy (42%) was picking up the occluding physical item(s) so that the digital item became visible. A second strategy (28%) involved dragging the digital item to a non-occluded tabletop zone, the physical item resting untouched. Moving the occluding physical item(s) instead of the digital one was a third strategy (27%). A final strategy (3%) consisted of moving both the physical and the digital items concurrently to remove occlusion. Participants leveraged the affordance of physical items to be moved not only on a flat surface but in all three dimensions.

In addition, they performed a considerable number of bimanual hybrid interac-

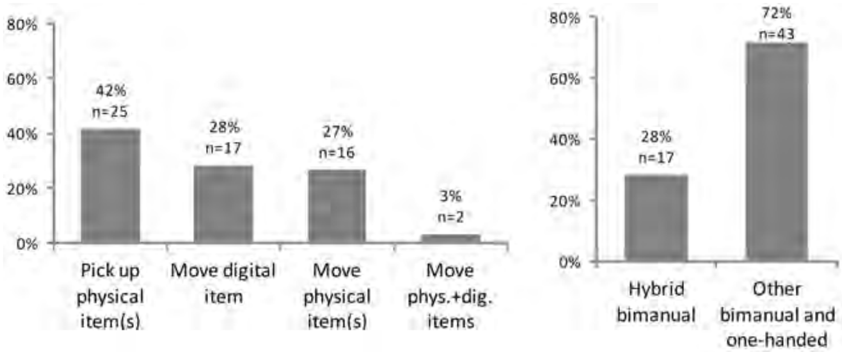

Figure 3.6: Strategies used for selecting occluded digital items [Steimle 2010b]

tions (28% of all instances). These involve manipulating physical items with one hand, while the other hand interacts with digital items. While we observed symmetric bimanual interactions, asymmetric interactions were more frequent. Participants typically picked up the physical item(s) with one hand (typically the non-dominant hand) and kept them in this hand while using the other hand to interact with the digital item (e.g., zooming). Finally, they placed the physical item(s) back on the surface.

The selection of one of the aforementioned strategies highly depended on the degree of occlusion. If the amount of occlusion was higher, users tended to pick up the physical item(s), whereas in cases of less occlusion, users preferred to move physical or digital items on the surface. This correlation between degree of occlusion and picking up physical item(s) was highly significant ($r = .53$, $p < .001$, $N = 60$). It was also highly significant for the degree to which the area surrounding the digital item was occluded by physical items ($r = .48$, $p < .001$, $N = 60$).

We also analyzed how the selection affects the overall hybrid arrangement of items. Groups should remain groups and should not be altered by a selection activity. We observed that the overall arrangement of items remained the same before and after the selection if the digital item was occluded by only one item or by several unordered items. Afterward, the occluding items could be easily placed back in the correct arrangement after selection.

Zooming

In general, participants frequently enlarged digital items in order to read the contents. If one or more physical items were close by, enlarging could lead to occlusion. We observed 23 instances of occlusion resulting from enlarging. On average, 30% of the item's surface was occluded (SD = 15%), with a maximum occlusion of 70%.

Despite the considerable amount of occlusion, participants did not find it to be problematic.

Since the tasks (for grouping and comparing) required access to specific information that was contained at similar positions on each document, participants could perform the pinch gesture to enlarge the document where the information was assumed to be located. While the document was being enlarged, the outer portions of the document expanded and moved under nearby physical items. The position of interest remained at the same non-occluded location and could easily be read. If the task required skimming or reading the entire document, occlusion turned out to be more problematic.

Enlarging items that were partially occluded at their outer edges resulted in a cluttered view. The physical arrangement did not clearly convey whether the occlusion was accidental or the items were purposely arranged in an overlapping manner to express a relationship. To avoid this ambiguity, in 65% of all instances participants immediately scaled down the item once they had read the information. Often enlarging, reading and scaling were performed in one integrated interaction using a single continuous pinch gesture. It is likely that zooming will be a less frequent activity in future tabletops, which will offer a higher resolution. Nevertheless, enlarging a view or an individual element is an important activity with many types of documents and visualizations, independent of the resolution available.

Moving

Moving items is a core interaction for managing a workspace. However, when using touch interaction to drag a digital item, one cannot move it across positions occluded by physical items. We were interested in how users would deal with this situation.

Figure 3.7 (left) depicts the three different strategies we observed and their frequencies. The most frequent strategy (48%) consisted of moving the digital item around the physical item. If this move-around approach would be tedious or even impossible, participants frequently leveraged hybrid bimanual interactions. In order to remove any obstacles in the motion path, they picked up the physical item with the non-dominant hand, moved the digital item with the dominant hand to its target position, and then placed the physical item back on the table. This strategy was used not only when a single item created the occlusion but also when entire piles had to be picked up. A third strategy involved first moving the physical item on the surface to free the path before dragging the digital item to its target position.

Piling

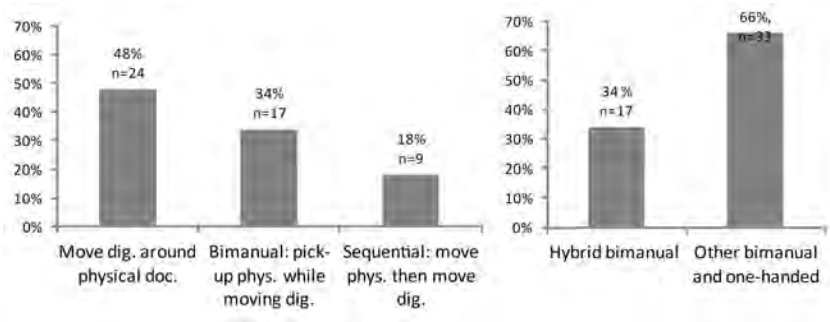

Figure 3.7: Strategies used for moving digital items over occluded positions on the tabletop surface [Steimle 2010b]

Participants' interactions with hybrid piles demonstrate that they intentionally generate occlusion by placing physical media items on top of digital media items. This spatial organization serves to represent and highlight a meaningful relation between the elements of a pile and thus rather than being problematic provides a visible representation of a conceptual relationship. As a consequence, it seems appropriate to distinguish between semantically meaningful forms of occlusion and forms that do not express semantic meaning. Occlusion in a hybrid pile can be semantically meaningful and desired by the user when creating the pile which we provided support in the previous chapter.

In addition to this positive semantic aspect of occlusion, problematic features were also evidenced. Once a hybrid pile was created, physical items often entirely covered the underlying digital items. As a consequence, there was no indication that there were digital items hidden under physical ones. In interviews, five participants reported this to be highly problematic. In order to indicate the presence of digital items, three participants slightly displaced the physical items of a hybrid pile so that the digital items remained visible.

3.2.2 Interim Summary

The results of the study revealed that users are willing to physically occlude digital contents in order to better manage their workspace. Moreover, when coping with occlusion, participants had highly effective strategies that rely on bimanual interactions. In general we found two main strategies for accessing occluded objects: moving or lifting physical documents, and dragging out digital items. The selection of strategies highly depended on the degree of occlusion as well as surrounding occlusion. Zooming of digital occlusion often resulted to occlusion and a cluttered

✓ Presence of physical documents hindered interaction with and display of digital ones
✓ Two strategies was found to access an occluded object: –moving or lifting physical objects –dragging out digital objects
✓ The degree of surrounding occlusion significantly influenced what strategies to employ
✓ Zooming digital documents often resulted in occlusion and a cluttered view
✓ Moving digital objects was found to be cumbersome and ineffective

Table 3.2: Summary of physical occlusion-related results from the hybrid media usage study

view. Regarding moving digital objects across the tabletop surface, our analysis underscore moving around physical objects as the most frequent approach. Table 3.2 summarize the results of this study.

3.2.3 Second Exploratory Study

In this exploratory study, our main objective was to gain insights about how people interact with digital objects together with various everyday physical objects – such as books, coffee cup, and laptop. We recruited six volunteer participants from our university department. One participant was female the rest male with a mean age of 28 years. All participants were familiar with interactive tabletops. In the study, we used a Samsung SUR40 tabletop running a Bing sample application, allowing users to search for images and geographical locations using the Bing Internet search engine. Query results included images and maps snippets with which users could interact using common multitouch gestures. The study involved single-user sessions that lasted approximately one hour.

We asked participants to organize a mid-size academic conference in the local city. This required participants to look for suitable hotels, a conference location, and a place for a social event. Participants were asked to physically document their results and draw a map of the locations using a pen and paper along with presenting their visual search results. In order to create a realistic hybrid setting, participants were encouraged to bring and place their everyday objects – such as cell phones, wallets, or coffee cups – on the display surface. Since there was no textual search provided in the sample application, an iPad was also supplied in case users needed to textually search for further information about hotels or locations. Following completion of the task, a semi-structured interview was conducted with each participant. Sessions were videotaped. We present the main results of this and previous studies and summarize as a set of requirements, presented next.

3.2.4 Requirements

Based on the results of both studies, we derived a set design requirements. Note that the results of the second study are categorized and integrated into the requirements. We also discuss previous work related to each of the requirements to see whether it has been addressed before.

R1. Support the positive aspects of occlusion

We frequently observed participants concealing digital objects for various purposes. For example, they did this to hide objects no longer being used following a focus or task switch, to reduce visual complexity of the workspace, or to group or pile a set of digital and physical objects. Participants (P2, P3, P5, P6) pointed out other situations in which they would want to deliberately hide an object, similar to minimizing a window in a WIMP interface. Similar to our findings from the previous study, participants also mentioned that techniques – such as automatic relocation of occluded objects – that do not adequately take useful functions of occlusion into account might be troublesome and decrease the feeling of workspace control.

Participants also confirmed the importance of accessibility and awareness of occluded objects and emphasized that visual clutter should be minimized. P3 added, “Having a mixed set of digital and physical objects overlapping one another does already [create] a lot of visual clutter. Adding a heavy system support [visualization] can make it worse!” This accords with the findings of Javed et al. [Javed 2011] and their suggestion to provide lightweight awareness techniques for tabletop applications.

Therefore, hybrid tabletop systems should indicate the presence of occluded items while respecting positive aspects of occlusion – as a concrete example, by displaying a halo around a physical item that fully occludes one or more digital documents. Moreover, selecting occluded items should be made easier to accomplish. Interfaces should permit users to temporarily display occluded items (perhaps as the result of hovering) at non-occluded positions to enable easy selection.

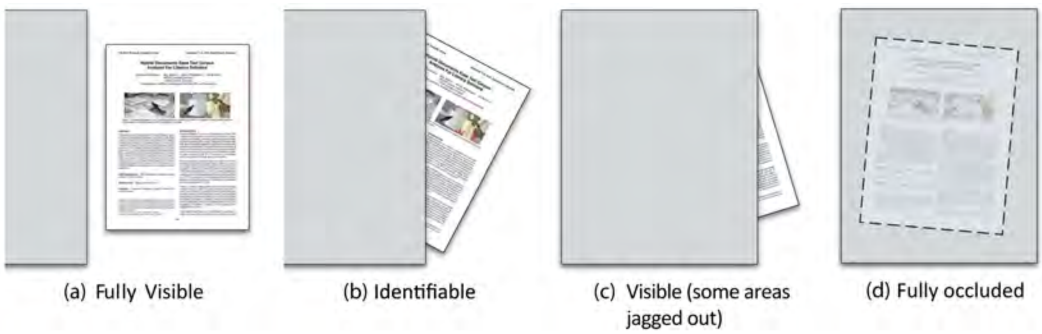

Figure 3.8: Occlusion continuous spectrum

R2. Support occlusion-aware workspace management

Our observations and participants' feedback revealed that to support effective work in hybrid environments requires not only providing awareness of hidden digital objects but also facilitating organizing and structuring the entire workspace. P3 showed moving a digital object through several physical objects and stated, "Moving this digital one in the hybrid setup is like playing a maze game. This is suboptimal and tedious in contrast to the fast digital move gesture." Ideally users should be able to relocate, pile, and structure virtual objects as easily as when there are no physical objects present on the tabletop surface.

R3. Reflect the continuous spectrum of occlusion

Once occlusion takes place, the severity of it can be assessed by measuring the covered-to-total-area ratio of the occluded item. We observed that occlusion occurs in varying degrees, ranging from full occlusion, with occluded objects not visible at all, to partial occlusion, to full visible. In addition, partially occluded elements can be identifiable or unidentifiable and accessible for manipulation or not. In general, we can classify different degrees of occlusion as a spectrum. An illustration of the different degrees can be seen in figure 3.8:

(a) All parts of the digital item are *fully visible*. There is no occlusion taking place. This degree is only stated for completeness.

(b) The occluded item is *identifiable* and interactable however some areas are covered.

(c) Some areas of the occluded item are jagged out so that the occlusion is still *visible* and the user can perceive that something is being occluded. In this level, the occluded item is mostly not easily accessible because of the high ratio of occlusion.

(d) The occluded item is *fully occluded*, and the user is unaware of it.

Occlusion-aware tabletop systems need to provide appropriate levels of information about occluded objects and ways to manipulate them. In addition, how occluded objects are represented should be sensitive to context and user task (see the next requirement).

R4. Reflect relevance of items

We observed that participants occasionally switch between items to complete their task. Items with higher *relevance* regarding the user's current context are more likely to be frequently selected. If the newly selected item is currently occluded, the user has to first resolve the occlusion in order to continue his task. This disruption can quickly become bothersome and can deteriorate structuring of the workspace. We infer that occluding relevant items to the user's current task might more bothersome. One solution when occlusions are cumbersome is to recognize the relevance of occluded items. There exist multiple primitive ways to approximate this property:

- Measuring how frequent an object is used. Occluded items with a high *usage frequency* likewise hold a high relevance.
- The current *field of view* or focal point of the user may give an indication on which items the user currently focuses on.
- *Spatial location of occlusion* on the tabletop surface can provide for a direct estimation of the item's relevance. This is only possible since users structure their workspaces based on distance to themselves.

Occlusion-aware technique should take the relevance factor into considerations and appropriately react to the occlusion.

R5. Support various resolving strategies

In the study we observed that participants employ three generic strategies to cope with occlusion. As illustrated in figure 3.9, these are:

(a) Relocating physical occluders (i.e., moving, sliding, or lifting the obstructing objects).

(b) Manipulating occluded items (i.e., dragging out digital items). This is only possible when at least some edges of the occluded item are visible.

(c) Applying a hybrid strategy (i.e., a combination of the above strategies performed as a bimanual interaction to resolve occlusion). Typically, it is composed of two actions concurrently: dragging out the desired digital item while relocating the physical occluder to free the space. Another hybrid strategy that we observed was peeling back a corner or side of a physical object (such as sheets of paper) while dragging out the underlying digital item.

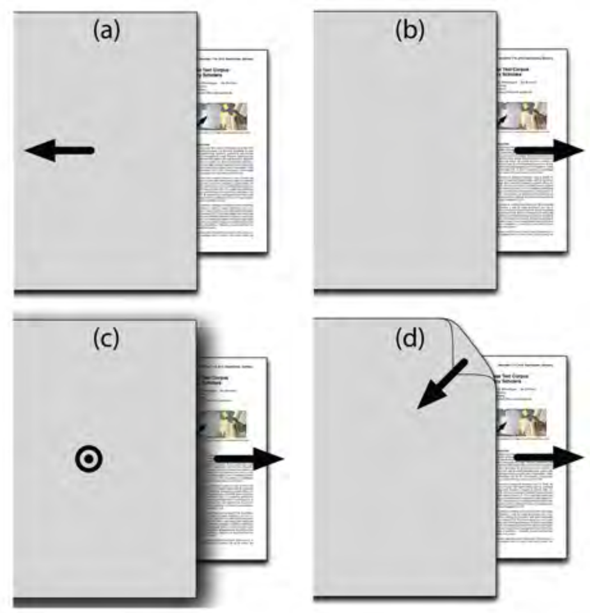

Figure 3.9: Generic strategies to resolve an occlusion: (a) relocating physical object, (b) dragging out occluded digital object, (c) and (d) hybrid bimanual strategies involving dragging out digital object and, at the same time, moving or lifting the physical occluder

Post-discussions with participants revealed that relocating physical objects to access the digital objects is the most natural way of coping with occlusion. The two other strategies turned out to be more cumbersome, and participants suggested that a potential system support should provide possibilities to alleviate dragging out digital objects or performing a hybrid strategy, particularly when digital objects are fully occluded.

R6. Support occlusion at different scopes

Discussion with participants revealed that system solution to cope with occlusion problems can be designed at different scopes. Occlusion can be considered *individually* for individual occluded items, *locally* for one occlusion group and *globally* over the whole table. Each scale has its own benefits and is better suited for particular tasks. For instance, in case of search tasks, there should be a possibility to resolve occlusion globally assisting the user in gaining a thorough overview.

R7. Provide physical anchors

Participants frequently expressed that in a hybrid setting physical objects are dominant. P2 said, "Physical objects are graspable and always atop the digital items and this makes manipulation of them much easier." Participants suggested explicit interactions – such as linking and moving multiple virtual objects using a physical object.

Table 3.3 provides a summary of the aforementioned requirements and the extent to which they are supported in the related work. We also show which of our them are met by previous work from section 2.1 and list our respective contribution of this thesis to overcome limitations of previous work. Our intent is not to mimic R1–R7 in our designs directly, but rather to use them as rationale for our concepts. Moreover, some of the requirements did not strike us as important observations until after we had designed and evaluated our techniques. In the next section, where our occlusion-aware concepts are presented, we will refer to each of the requirement that is supported by our designs.

	Requirements	Supported by previous work?	Contribution of ObjecTop
R1	Support the positive aspects of occlusion	○	ObjecTop interface allows for occluding digital objects
R1a	Support easy and quick access to occluded items	◐	ObjecTop provides support for awareness of and accessing to occluded items
R1b	Minimize visual clutter of workspace	◐	Halo visualization introduces less visual clutter
R2	Support occlusion-aware workspace management	○	ObjecTop offers a set of organization-supporting techniques
R3	Reflect the continuous spectrum of occlusion	◐	The gradual access technique enables semantically zooming into occluded objects before full access
R4	Reflect relevance of items	○	ObjecTop represents various levels of awareness about occluded objects based on their proximity to the user
R5	Support various resolving strategies	○	ObjecTop facilitates both dragging out and hybrid strategies
R6	Support occlusion at different scopes	○	Access-supporting techniques assist resolving occlusion in various scopes
R7	Provide physical anchors	◐	Hybrid binding gestures allows for persistently binding digital to physical objects

Table 3.3: Overview of our requirements for occlusion-aware tabletop interfaces and the extend to which they are covered in the previous work. ○and ◐show whether the requirements have partially been addressed in the previous work or not, respectively.

3.3 ObjecTop: An Occlusion Management Framework

In response to the key requirements established in section 3.2, we propose *ObjecTop* as an occlusion management framework for hybrid tabletop systems. It consists of a set of visualizations and interaction concepts supporting knowledge workers in such settings for the following basic practices:

- obtaining awareness about and overview of digital occluded items
- accessing and selecting occluded digital items
- interacting with digital objects across physical objects

We first present the general visualization concepts that provide awareness about occluded items, particularly including *proxy* representation of occluded items. We then describe two access-supporting techniques – namely, *DragView* and *Pressview*. Finally, we present the organization-supporting techniques, which include *teleport*, *remote zooming*, *hybrid piling or hiding*, and *binding*.

3.3.1 Awareness-Supporting Concepts

Recalling R1, we design our occlusion-aware techniques in a way that they support the duality of occlusion. This means that our solution considers and allows for employing both the positive as well as negative aspects of occlusion. The ObjecTop system allows users to intentionally occlude digital items for the purpose of managing space or deliberately hiding items. Thus, once occlusion takes place, we do not automatically relocate the digital objects. Instead, the system informs users by showing unobtrusive visualizations of occlusion and representing occluded items through *proxies*.

Proxy

Informed by the literature analysis (3.1), we employ a proxy-like representation of occluded items (cf. figure 3.11) which is basically a semantically scaled representation of the actual item. Proxies serve two main purposes – namely, delivering an adequate level of awareness to users about occluded digital items and providing an additional strategy to resolve occlusion (see access-supporting concepts). Typically, they are placed on the nearest edge of the physical occluder and therefore the spatial relationship between occluded items and physical occluder is preserved. As an example, figure 3.10 shows that the media player is occluded and its proxy is visualized right on top of the occluder edge.

However, in situations where there is not enough space to place the proxy right on the nearest edge of occluder (due to, for instance, surrounding occlusion or tabletops' edges) it is placed far from the actual object. Figure 3.10 illustrates an example that due to the lack of enough space near to the edge of tabletop, the proxy is located far from the occluded image object. To mitigate this problem, we designed a *beam* that connects the proxy with its actual occluded item to convey some spatial information about occlusion. The direction of the line indicates the orientation to which the occluded item is placed. The width of the line also gives a rough estimation of how far the proxy and the actual object are from each other. The beam helps users to approximate where an occluded object is located under a physical occluder and could potentially support users in moving the occluder to resolve occlusion (R5).

Figure 3.10: Proxy placement – beam connects the actual object and its proxy representation (Highlighted objects indicated by dashed lines on the book and document are not part of the interface and are added in Photoshop for ease of understanding.)

Drawing upon R3, proxies can represent different levels of awareness about the occluded object. It ranges from showing no information about the hidden objects to uncovering the original object with various intermediate levels in between. We represent this spectrum by three main levels as depicted in figure 3.11:

Presence provides presence information about an occluded item under the physical object. This includes the type, number and rough location of occluded objects. This level minimizes visual clutter, providing only an indication of the presence of occluded objects.

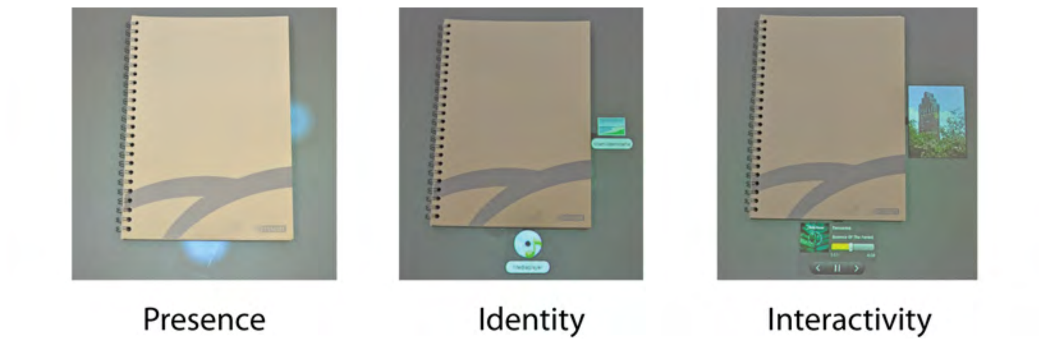

Figure 3.11: Different levels of detail represented by the proxy

Identity showing information about the type and identity of the occluded objects – such as the icon and the name – so that users can identify the occluded object.

Interactivity shows detailed preview information – such as a thumbnail or miniature version of the occluded object. Similar to the thumbnail view on the Windows Taskbar, it provides basic functions so that users can interact with an occluded object without retrieving it. Although this representation occupies more space than the other awareness levels, it supports a high level of engagement with the occluded objects.

Full access the occluded object is detached from the halo and moved to a position outside the area of the occluder.

The selection of varying levels of detail depends on the importance or relevance (R4) of the occluded objects to the context of the user. Since tabletops enable a spatial arrangement of digital and physical objects on their surface, the degree of importance of occluded objects is determined by their proximity to the user. Based on the in-depth analysis of Sellen and Harper on physical workspaces [Sellen 2003] as well as findings from the study on reachability for tabletops [Toney 2006], we feature three functional zones based on the distance to the user:

Active nearest zone to the user within the arm reach. Objects placed on this zone (physical or digital) are highly related to the context of users and frequently used. In this zone, when a physical object occludes a digital one, the highest level of awareness (i. e. the interactivity level) is shown by the proxy. Thus, the user can quickly find and engage with the occluded objects that are highly related to her context.

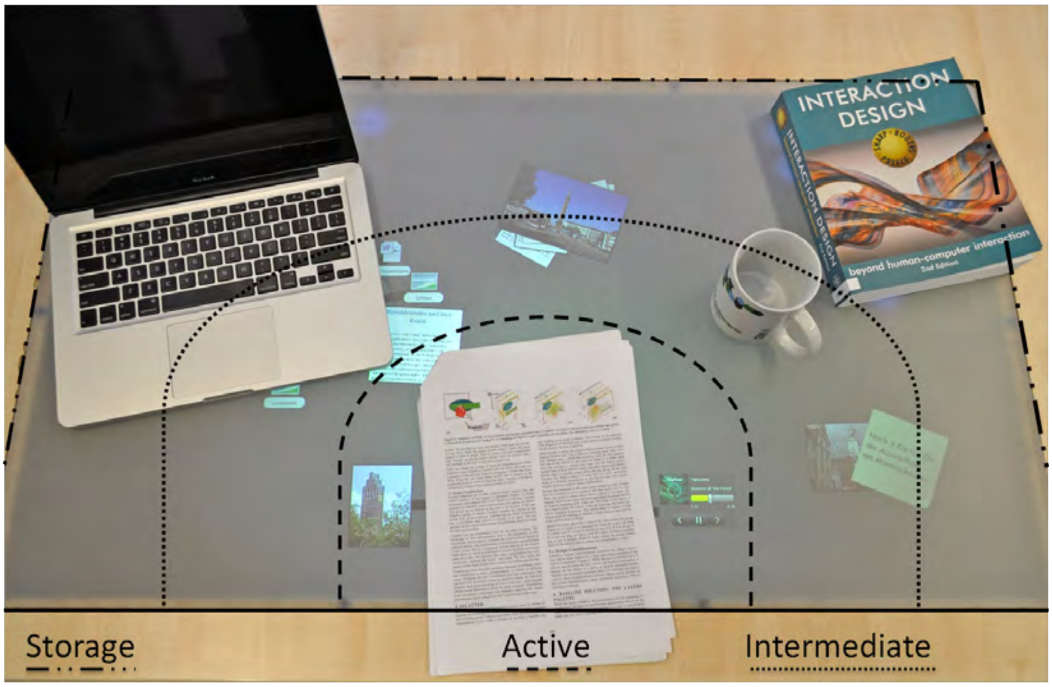

Figure 3.12: Proxy representations on each functional zone

Intermediate is a zone at arm's reach for temporarily placing objects that are not needed. Once occlusion occurs, the proxy provides awareness such that users can easily identify digital objects

Storage this is the farthest area on tabletop where the user can reach objects by leaning over the tabletop and is used for long-term storage of objects. In the case of occlusion, only a halo is represented indicating the presence of occluded objects.

These zones are illustrated in 3.12. By mapping different levels of detail to the different zones, the system offers a trade-off across the tabletop zones between user's intention as well as introducing visual clutter by the proxy representation. This means that although the interactivity level introduces more visual clutter in the active zone, it facilitates user to effectively interact with the frequently used occluded digital and physical objects. For the objects occluded farthest away and less frequently needed by the user, the system provides low-level representation thus, less visual clutter is introduced.

Displaying the basic proxy representation enables design of access- and organization-supporting techniques discussed in the next sections.

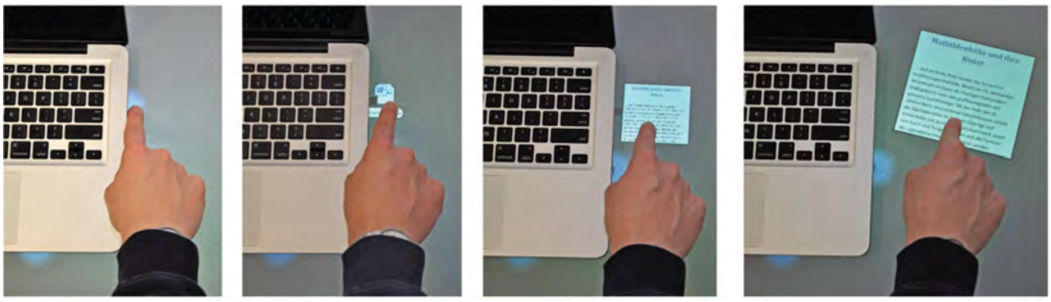

Figure 3.13: Accessing an occluded document by dragging out the proxy

3.3.2 Access-Supporting Concepts

In this section, we present interaction concepts and techniques to support accessing occluded digital items. These techniques are particularly designed to support resolving strategies identified in our study (R5). We first present DragView, a technique to support dragging out digital items, followed by the PressView technique (i.e., a pressure-based physical interaction helping users to bimanually resolve occlusion). Moreover, each of the techniques provides consistent and intuitive ways for resolving occlusion at various tabletop scopes (R6).

DragView

As mentioned above, proxies not only provide primitive awareness of occluded objects but also enable access to them. In our design we leverage the natural metaphor of dragging out to enable a peak under the digital items covered by a physical occluder. Our design provides *gradual access* to occluded objects. While the user is dragging an individual proxy, interactive previews of increasing levels of detail are shown before eventually displaying the object itself in full detail. Inspired by semantic zooming notions, the four levels of detail described above are supported (R4): presence, identity, interactivity, and full access. If a user releases an object before it is moved to full access, the object automatically snaps back to the presence representation. The gradual access technique is illustrated in figure 3.13.

Our framework differentiates from the one presented by Javed et al. [Javed 2011] mainly in two main respects. First, in their work, each occlusion level is supported by one or more techniques individually. Those techniques that support either presence or identification level (for instance, Glow, Icon, Pile, Minimap, shown in figure 3.4) do not support interactivity or access to occluded objects.

Second, other techniques such as Move and Replicate described in their work, enable full access to the occluded objects by always moving them to non-occluded areas. As already noted in the related work section, this approach consumes a great amount of space and introduces high amount of visual clutter. In contrast, our proxy representation provides presence and identity while minimizing visual clutter and, at the same time, allows users to gradually access various levels of detail of occluded objects. Moreover, the gradual access offers a lightweight way to peek at an occluded item or an entire occlusion group at the desired level of detail.

PressView

Based on our study findings, we design an interaction technique called *PressView* that supports a hybrid strategy to access occluded objects. In response to R5, it leverages and combines both physical manipulation of occluders and the digital drag-out technique. Inspired by *wringing out* as a metaphor, the user can press down the physical occluder lightly such that the proxies representing the underlying objects would temporary squeeze out (i.e., showing more levels of detail). This technique also supports resolving occlusion in different scopes (R6).

When the PressView is applied on a single side of the occluder, only those proxies will show up that are positioned near to that side. In order to get an overall preview of all occluded objects, users can perform the PressView roughly around the center of the occluder. Pressing on an empty space on the tabletop gives users a global overview of all occluded objects. While the PressView is active, users can drag one or multiple proxies out in order to directly access the original object, thus supporting a bimanual resolve of occlusion (R5). The applied force by the user can also be mapped to the degree of awareness represented by the proxy (i.e., as the user presses more the proxy more information about occluded items).

The PressView technique offers three main advantages. First, it does not require any graphical widgets to be displayed on the tabletop nor any complex gesture for triggering it. Second, it easily supports resolving occlusion in situations when moving the physical occluder is barely possible, either because of heavy or complex occluders (such as a laptop or a thick book) or due to a specific arrangement of multiple occluders (such as a number of sorted paper documents).

Third, in 2D desktop environments one of the common activities is to find and retrieve a specific window amongst all hidden overlapped windows. In order to do so, users need first to obtain an overview of all open windows and then select the one they searched for. With current window management systems this can be efficiently achieved by tiling (Expose mode in Mac OS) or cascading (like Windows Flip 3D) all open windows. The PressView technique resembles and supports the same behavior

(a) PressView on the side

(b) PressView on the center

(c) Global PressView

Figure 3.14: The PressView technique resolves occlusion at various scopes on the tabletop

in hybrid tabletop settings by enabling users to get a global overview of all occluded objects.

3.3.3 Organization-Supporting Concepts

Organizing information and objects on the physical desktop is a semantically valuable and common activity [Mander 1992]. Knowledge workers frequently organize information in form of piles or loose groups on their workspace. While such activities are well-supported in only physical or only digital settings, our interaction concepts presented in this section aim to support such frequent organizational activities in hybrid tabletop settings. The key principle used in the design of our technique is based on Guiard's seminal work [Guiard 1987] in which he introduced the Kine-

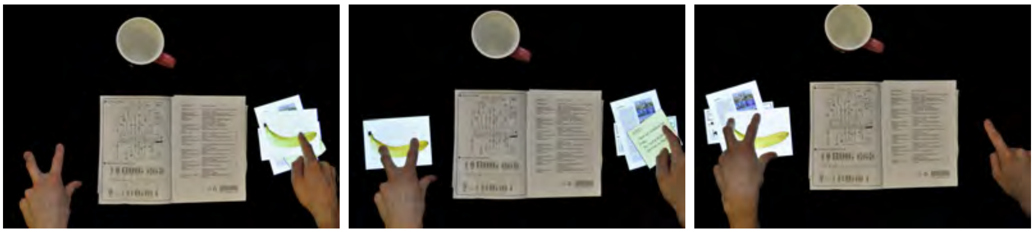

Figure 3.15: Teleport gesture – tripod finger gesture indicates the destination, and digital documents are flicked toward the destination using the dominant hand

matic chain model for asymmetric bimanual interaction. According to this model, the roles of the two hands are such that the non-dominant hand sets the frame of reference in which the dominant hand operates.

We take advantage of two-handed interaction and design a set of multitouch gestures in which the non-dominant hand indicates a destination or selects an object while the dominant hand performs a specific gesture to accomplish a task. The bimanualism incorporated in the design of organizational techniques increases the reachability of users across all tabletop zones. In the following, we in turn describe three organizational techniques that facilitate hybrid moving, zooming, piling, and binding.

Teleport

Moving is one of the most frequent activities when dealing with either physical or digital objects. In hybrid settings, however, moving of digital objects can become challenging due to the presence of physical objects on tabletops. Physical objects may block the path across which a virtual object can be dragged. Users need to either move the physical object out of the away or move along a less direct path. We designed a bimanual moving technique called *Teleport*: one hand indicates the destination with a tripod-like, three-finger gesture, and the other hand flicks the object toward the finger tripod (R2). This interaction is depicted in figure 3.15.

In comparison with free-flicking gestures, Teleport is more precise by specifically indicating the destination. This is critical in hybrid settings, as otherwise flicked objects could end up covered by an occluder. In addition, it can be employed to quickly collect many digital objects in a pile, by flicking multiple digital objects from across the tabletop to the same destination.

Remote Zooming

If several physical objects cover a tabletop area, there might not be sufficient uncovered adjacent display space available for gradually accessing occluded items using drag-out or pressure-based accessing techniques. We designed a bimanual gesture that enables users to remotely inspect occluded objects on any empty spot on the tabletop surface (R2). While touching an individual proxy with one hand, the other hand performs a pinch gesture on any other area of the tabletop display as depicted in figure 3.16. The temporary view is placed slightly off the center of the pinch gesture to ensure it is not obstructed by the user's hand and fingers. Users can also inspect multiple items in series by sequentially referencing on them using their dominant hand.

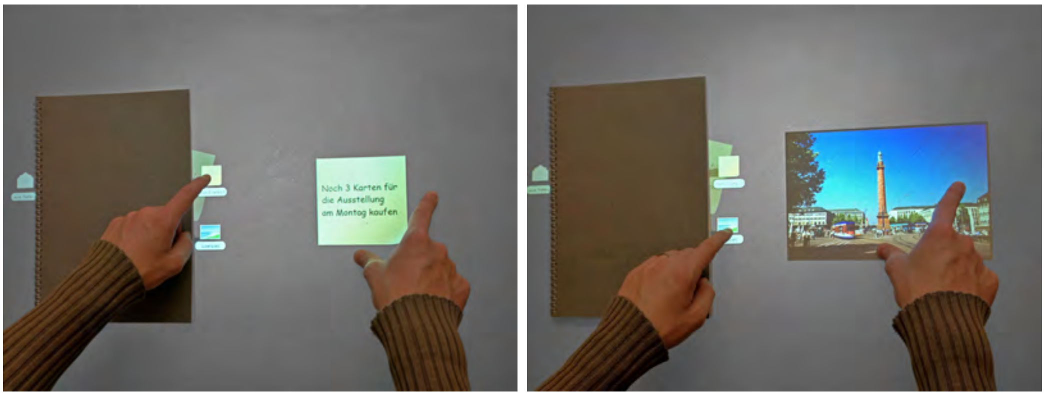

Figure 3.16: Remote zooming gesture – the user first indicates the proxy and then performs a pinch gesture to zoom into the document on an empty space of tabletop

Hybrid Piling

Another common workspace organizing activity is piling or stacking to create groups of physical and/or digital objects [Mander 1992] or to reduce visual clutter by deliberately hiding objects underneath other objects. To support this activity, users can perform a simplified variant of the Teleport gesture by placing one or two fingers on the halo (or close to the edge of the physical occluder) and performing a flicking gesture. The technique is illustrated in figure 3.17. In this way users can create hybrid piles or deliberately hide one or multiple virtual objects underneath a physical one (R1).

Hybrid Binding

People frequently bind or associate physical objects to one another, for instance, by stapling multiple sheets of paper together or by sticking Post-it notes onto docu-

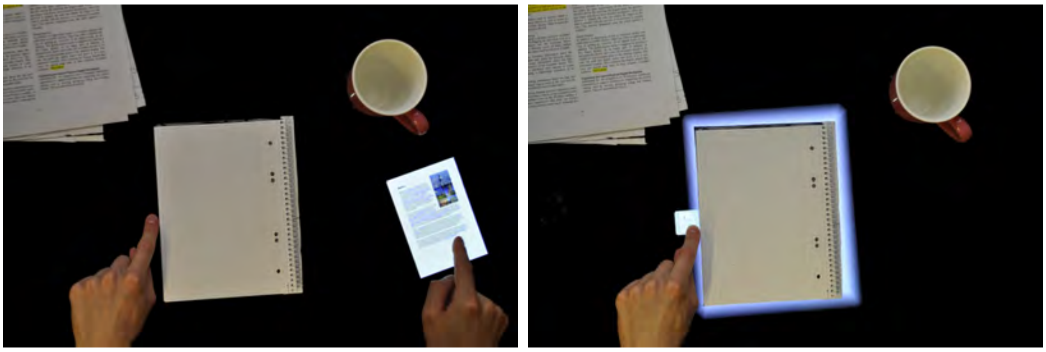

Figure 3.17: Hybrid piling gesture – the non-dominant hand indicates the physical object under which the dominant hand flicks the digital object

ments. Binding indicates some relation between the attached objects and can enable easier access and comparison once objects are bound. We designed a bimanual gesture so that users can attach digital to physical objects (R7). To do so, the user touches the digital item to be bound and performs a swipe gesture along an edge of a physical object (cf. figure 3.18). The swipe gesture is inspired by the action performed to close a zipper. While swiping, users can adjust the size of the area used to dock the digital object to the physical one. Once attached the digital item automatically moves with the physical object when it is moved. To detach, users can simply drag off the docked digital object.

Table 3.4 summarizes the three main interaction concepts of ObjecTop along with interaction techniques.

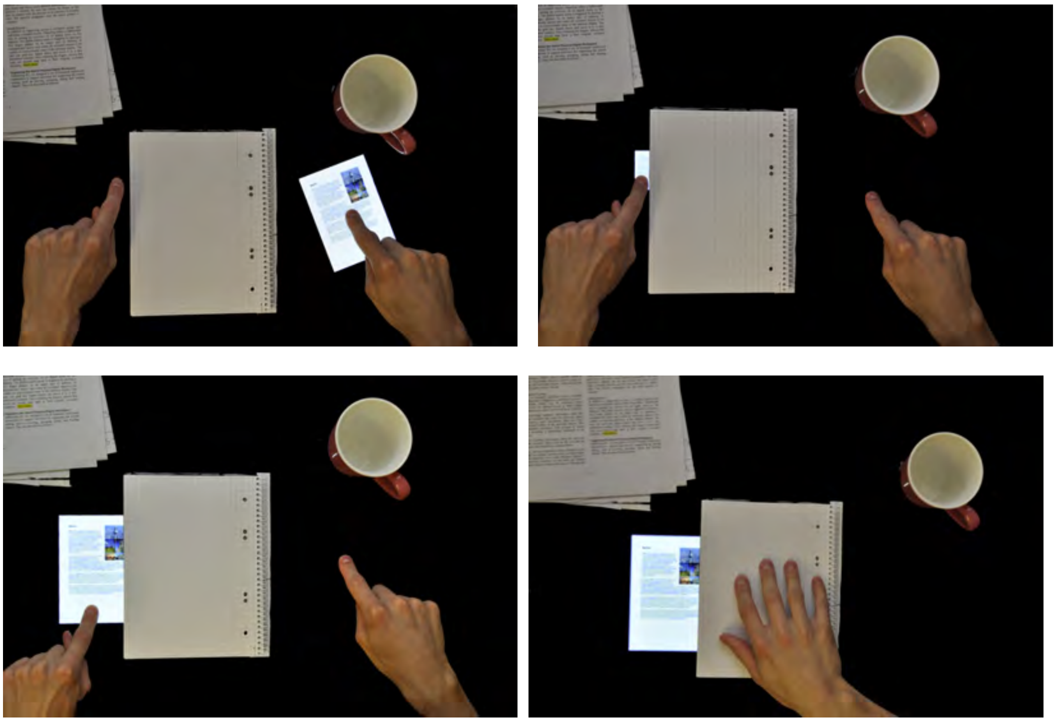

Figure 3.18: Hybrid binding gesture – the user first selects the digital object and then performs the swipe gesture along the edge of the physical object for a permanent hybrid bind between two objects

Awareness-Supporting Concepts		
Name	Purpose	Description
Proxy	Representation of and enable access to occluded objects	Displays various levels of detail: (i) visibility in storage areas (ii) identification in intermediate areas (iii) interactivity in active areas
Access-Supporting Concepts		
Name	Purpose	Description
DragView	Gradual access based on dragging-out metaphor	Dragging out the proxy to see more levels of detail, and eventually access the occluded object
PressView	Gradual access based on wringing-out metaphor	Pressing on a physical object to see more details of underlying digital objects and drag it out with the other hand (bimanual approach)
Organization-Supporting Concepts		
Name	Purpose	Description
Teleport	Targeted flick of digital objects across physical ones	Finger tripod as destination and flick the object toward it
Remote Zooming	Inspecting an occluded object on any empty spot of tabletop	Select an object and perform pinch on non-occluded area
Hybrid Piling	Pile or hide digital objects under physical ones	Select a physical object and then flick digital ones over
Hybrid Binding	Permanent linking digital to physical objects	Select a digital object and perform a swipe gesture along an edge of the physical object

Table 3.4: Summary of occlusion-aware interaction concepts presented in the ObjecTop interface

3.4 Implementation

We implemented the aforementioned interaction techniques as a functional prototype called *ObjecTop* running on the tabletop system described in section 2.2.1. It has a display size of 100 x 60 cm and an additional outer frame with a width of 30 cm on both sides. The tabletop is a back illuminated with diffused infrared light. The integrated camera is a Point Grey Flea2 CCD Camera [1], with a maximum resolution of 640 x 480 pixels at 80 frames per second. The integrated projector supports 1920 x 1080 pixels and has a resolution of roughly 48 ppi, which is sufficient for reading documents on tabletop. 3.19 shows a photo of our tabletop system.

(a) Overview (b) Inner housing

Figure 3.19: The tabletop system of ObjecTop

We used *reacTIVision* [2] as our underlying touch and fiducial marker tracking framework. It is an open source, cross-platform computer vision framework for the fast and robust tracking of fiducial markers attached onto physical objects, as well as for multitouch finger tracking. We augmented physical objects with the proposed amoeba fiducial markers from [Jordà 2007] to identify and track them on the tabletop surface. Due to the low resolution of the tabletop camera, we had to increase the marker size. Depending on the footprint of each object, we attached fiducial markers with the size of from 8 to 12 cm. They, however, perfectly fit underneath the objects and allow for tracking.

In order to realize PressView technique, we modified our custom-built tabletop system to sense pressure applied on its surface. For that purpose, we used Phidgets [3] force resistive sensors (cf. figure 3.20a), which are paper-thin and can be easily

[1] http://www.ptgrey.com/products/flea2/flea2_firewire_camera.asp
[2] http://reactivision.sourceforge.net/
[3] http://www.phidgets.com/products.php?category=33&product_id=1131_0

placed between the edge and protection layer of the tabletop as shown in figure 3.20b. The force sensors can measure force between almost any two surfaces up to 2 kg. To ensure the most accuracy, we attached a small (blue) pad directly on the sensing disc. This will ensure that all the force is applied directly to the sensing disc and not to the surrounding surface.

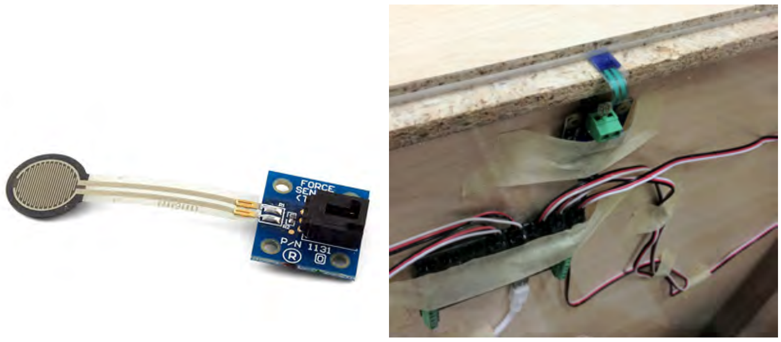

(a) Force resistive sensors (b) Embedded in the tabletop

(c) Inner view of our tabletop equipped with ten force sensors

Figure 3.20: Augmenting our tabletop system with pressure sensors

To recognize a press-down event on physical objects placed on the tabletop surface, we embedded ten equidistant of those force sensing resistors around the tabletop (cf. figure 3.20c). Two USB interfaces on each side connect respectively five force sensors to the tabletop's computer. When pressure is applied on the tabletop surface or a physical object, the force is distributed along all sensors accordingly. We use this physical principle to calculate the pressure point. The fraction of measured force of each sensor to the total force is used to approximate the surrounding sensors to the pressure point. Given the known positions of sensors around the tabletop surface the system computes the pressure point using vector spaces as follows:

$$\begin{aligned}
force_i &:= max(0, sensor_i - offset_i) \;\; \forall i \in \{1, ..., n\} \\
forceSum &:= \sum_i force_i \\
forceRelative_i &:= \frac{force_i}{forceSum} \;\; \forall i \in \{1, ..., n\} \\
pressurePos &:= \sum_i forceRelative_i \cdot sensorPosition_i \\
&\text{Where } sensor_i \in \mathbb{R}, sensorPosition_i \in \mathbb{R}^2
\end{aligned} \tag{3.1}$$

This approach requires a pre-calibration of the pressure sensors to exclude the weight of already placed objects on the surface. Our system had to be calibrated when new objects are placed on or removed from the table. The gained calibration is saved in the $offset_i$ variables inside the equations above. When the distribution of the force varies without significant changes in the summed force, the system can infer that one or more objects were relocated and performs the calibration automatically.

In this way, the system approximates the position of a pressure event in an area of 7 x 7 cm, sufficiently precise to recognize pressing down events on most of everyday objects. To improve the pressure accuracy, recent advances in pressure sensing technology [Leitner 2011, Davidson 2008] open up more promising and accurate approaches that can be considered in future studies. The system prototype described above enabled us to evaluate our interaction techniques presented in the next section.

3.5 Evaluation

We evaluate our interaction concepts and techniques through two user studies conducted in a realistic hybrid setting. With the aim of the first study, we want to validate our main design decisions and gain insights on how people use our techniques to cope with the problems of occlusion. Based on user's feedback, we revisit some of our techniques and identify a fair baseline as well as several key aspects of our occlusion management framework to be examined in a formal comparative experiment. We then conduct a second study in which we quantitatively examine our interaction techniques. We report the studies below in turn.

3.5.1 Study1: Exploratory Experiment

The main objective of this study, exploratory in nature, was to evaluate our design decisions. We were particularly interested in observing the users' response to our basic techniques. Particularly, we tested the *usefulness* and *viability* of the main features of our occlusion management framework – namely, proxy representation, functional zones, and access-supporting techniques. In order to avoid overwhelming users by all features of the occlusion management framework, we intentionally excluded organization-supporting techniques for this study. Another objective of this study was to find several main aspects of the system that could be evaluated in a formal controlled experiment. We chose realistic meaningful tasks to be accomplished with a set of everyday objects resembling a true hybrid knowledge work setting.

3.5.1.1 Method

Participants
We recruited eight volunteer participants (two females - six males) from which two were left handed. They ranged from 17 to 35 years old. Seven participants were experienced knowledge worker and one a preparatory school student. All of them had experience with multitouch devices – such as smartphones or trackpads – and no experience with interactive tabletops. The lack of experience with interactive tabletops ensured that they did not adapt any behavior or patterns to compensate for occlusion on such devices.

Study Setup
The study was conducted in our lab setting. We used the prototype system presented in 3.4 for this study. Our prototype offers a workspace of 130 x 105 cm with a display size of 100 x 60 cm leaving 30 cm rims around the display surface. This is representative of the space available on a typical desk. The rear projection had a full HD resolution of 1920 x 1080 pixels. The participants could interact with digital documents using typical multitouch gestures for moving, rotating, and zooming. Physical objects could be placed and manipulated on the display surface although, putting on the surrounding non-display areas (rims) of the table and stacking on top of each other were not allowed. The setting is shown in figure 3.21.

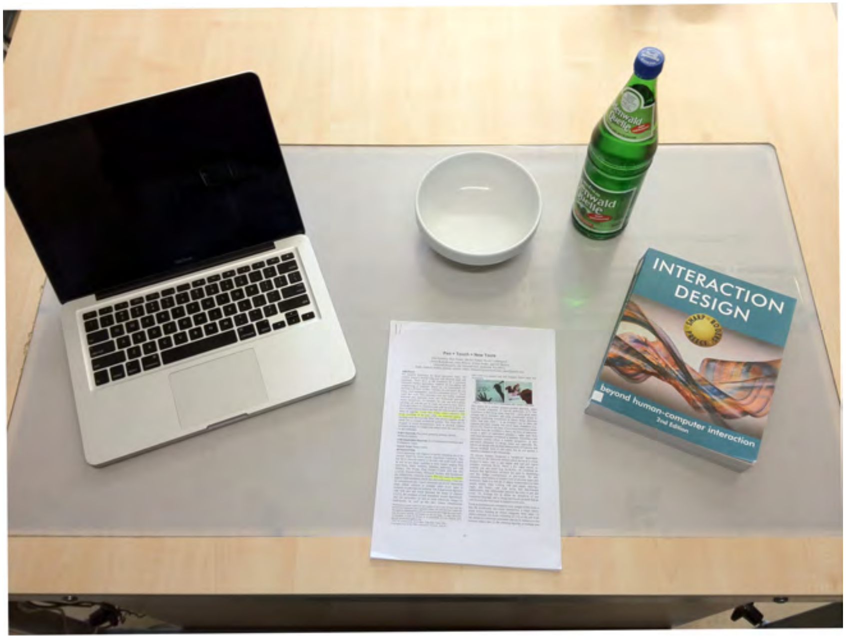

Figure 3.21: Study setup and initial arrangement of physical objects before each task

Materials

We chose a diverse set of everyday physical and digital objects to emulate a realistic setting and to create natural behavioral patterns. The information space used in the study was about three well-known sights in the local city. Physical objects used in the study included a thick book about sightseeing in the local city (roughly A5 in size and weighed 1 kg), a flyer about a particular sight printed on a 20 cm square cardboard (representing a typical lightweight object), a white ceramic bowl (16 x 16 x 8 cm and weighed 300 g), a water bottle that was 1 / 3 filled, and a 15¨ laptop weighing 2-3 kg with the lid open. The laptop showed questions and provided space for answers. All physical objects were recognized using fiducial markers. We used the following digital objects: seven text documents containing information about the three sights in the local city, six images about the sights in both landscape and portrait modes, one email containing sender, subject and body, and one Post-it note.

Tasks and Procedure

In our task design participants were asked to perform search and grouping tasks where all needed information was distributed over multiple physical and digital objects. In order to familiarize participants with the hybrid information space, we started with less clutter and a simple grouping task. As the participants progressed, we increased the level of clutter as well as the complexity of the task. Additionally, each task required the user to work with printed and digital media concurrently. Before the start of each task we set objects to the initial arrangement shown in figure 3.21 to control the starting conditions.

Each session was begun by a short briefing about the system to allow participants to manipulate digital objects using multitouch gestures. We then added the physical objects to the information space. Then we introduced the proxy visualization and both dragging out and PressView interaction techniques. The first task was to perform a basic grouping based on type of digital objects. In the next task, the participants were asked to find the relevant objects about only two sights and group them into two separate piles. In this case, the objects related to the third city sight served as distractors.

In the next task, four detailed questions about the sights were presented to participants on the laptop screen, where also the answers had to be typed in. The questions required a more thorough and detail search of the information space. In the last task, we asked a relative complex question about one of the three city sights, for which participants had to revisit nearly all objects. This time the answer had to be written down on stickies attached to a page of the book. Finally, we conducted a semi-structured interview with participants.

3.5.1.2 Results

We reviewed 11 hours of video material and created transcripts for each participant. Each transcript included reports and timing information of relevant behavior – such as handling physical objects, resolving strategies used to tackle occlusion, and utilization of interaction techniques. Feedback as well as suggestions on how to cope with occlusion in general were also added. All transcripts were finally summarized to gain insights into different behavioral patterns. The resulting data supported extracting generalized patterns about the main features of our system, presented next. We also supplement some findings with statements from the interviews.

Resolving Strategies

Our study design allowed participants to employ both generic strategies (relocating physical occluders or using our access-supporting techniques) to resolve occlusion. We therefore analyzed under what circumstances which strategy was used. We observed that physical objects – such as those that cast 3D occlusion and are heavy or difficult to handle – have been rarely relocated by participants.

We also observed that highly complex objects – such as the laptop and book – that are heavy and occlude larger area on the tabletop became sort of *stationary occluders* which have not been replaced through the study. For these objects, our analysis revealed that in 78% of occlusion resolve cases participants used our techniques. On the contrary, it was found that the paper flayer and bowl were more often moved or lifted. For these objects lifting or moving to access the underlying digital object was preferred by all participants.

We, however, observed that due to the lack of a group access in our design, participants had to drag out and inspect proxies individually. This was particularly bothersome when the number of proxies were greater than three and, eventually, resulted in either lifting or moving away the physical occluder.

Proxy Representation and Placement

Nearly all participants appreciated the idea of proxy representation. They found proxies practical and useful as a mean to access *fully* occluded objects. However, some participants had difficulties distinguishing actual digital items from the proxy representations. This problem primarily raised with the miniature representation as it looks very much like the actual occluded object. Most of the participants stated that proxies need to be visually more distinguishable than their original items. P2 stated, "I cannot always say in an instant what is a proxy and what not".

We observed that in cases where a sheer number of proxies were resting around an occluder, participants had difficulties completing the task effectively. Low level of detail as a representation alleviate part of the difficulties but nevertheless participants asked for *group-based visualizations* instead of individual proxies representation. They postulated that with more than three to four occluded objects (resulting in three to four proxies resting aside the occluder), the system should switch to a group visualization showing only a lightweight visualization around the occluder. Three participants (P2,P3,P5) suggested visualizing occlusion for hybrid piles always with such group visualizations to reduce visual clutter in the storage (peripheral) zone.

Placement of proxies was extensively discussed with participants. They reinforced the importance of spatial relationship in placing proxies. Nearly all of participants wanted the proxies nearer to the occluded object even at the cost of decreasing the level of awareness. A statement about this problem was if one sees occlusion not only as something bad but also as a form of organization, then something has a reason to be occluded and therefore it makes sense to position it nearby (and not far away because of its size. In other situations where placement of proxies became crucial, a participant (P3) reported, “I can use the spatial relationship as a cue to effectively move an occluder.”

Dragging-out Access

All participants regarded the drag-out access as a reasonable solution to the problem of free space in hybrid environments. P6 stated, “I can guess what’s inside but it should not take that much real estate on the screen”. With respect to each level of detail of the proxy representation, we received feedback from the participants which we list as follows:

Presence All participants stated that this level of detail clearly offers the least visual distraction. Some participants critiqued that they had to always interact with proxy to be able to identify what is occluded. This also matches the observation in which participants accessed proxies multiple times in order to identify the occluded objects. We also found that participants used the pressure-based technique to alleviate these problems. Our video analysis also confirmed this fact that *visibility + pressure-based access* turned out to be a frequently used combination by almost all participants.

Identity icon was perceived to be the most suitable representation by almost all participants. They provide necessary information about the type of the occluded objects. One participant mentioned *Icon delivers the right amount of information.*

Interactivity miniature helped participants to extract more information from the proxies than other representations. Yet it was interesting that the participants felt overwhelmed with too much visual clutter introduced by this type of visualization. They stated that particularly the working area of the tabletop become very cluttered with the miniature representation.

Pressure-based Access

Participants all appreciated the pressure-based access as they offer very lightweight and intuitive interaction for resolving occlusion. We observed that they frequently used PressView technique for quickly identifying occluded items. P7 noted, “I really like squeezing out proxies. It’s [a] clear metaphor that one can quickly understand.” A few participant used PressView at the end of each task to ensure that they browsed all necessary objects on the whole workspace. P8 stated that “by pressing on an empty area of the surface," they could quickly ensure that no objects were left out. The discussion with participants revealed that the PressView could become tiring particularly when they had to hold pressure on a physical object while inspecting squeezed out proxies. They suggested applying some histerisis or animation (fading out) could mitigate this problem.

Functional Zones

All of the participants noticed the different representations across the tabletop zones but only two could foresaw its reason (i.e. due to functional zones). The rest of participants mentioned other reasons for the chosen representation level – such as the type of occluded digital object, space left on the tabletop, usage frequency, or importance of occluded items. We believe size of tabletop used in the study was relatively small in order to have different representations across its zones. In many cases, the footprint of occluders was shared in two different functional zones, which resulted in showing different levels of details for one occlusion group. This turned out to be very confusing for the participants. Therefore, they suggested other factors to increase the level of detail of occluded digital items that:

- are frequently used.
- relate to the currently used object. Backing up the semantic relationship.
- are associated to the current task in hand / share the same context.

Nearly all participants preferred visibility as the default level of awareness for all tabletop zones. Interactions – such as the PressView technique – would serve as a quick means to identify items. The participants stated that they prefer a combination of visibility and group-based visualizations. The main findings of the study are summarized in table 3.5.

3.5.2 Design Improvements

Based on our findings, we propose a set of design improvements for our techniques, which are detailed here. We moreover improve our system prototype and implementation described at the end of this section.

3.5.2.1 Halo as a Group Visualization

We improved our design so that once occlusion takes place, our system visualizes an interactive halo representation to provide awareness of occluded items. The

✓	Replacing lightweight physical occluders was easier than dragging out access.
✓	For complex occluders, drag out strategy was mostly used to resolve occlusion.
✓	Proxies were found to be helpful particularly for accessing fully occluded items.
✓	Icon representation turned out to the most suitable level of awareness.
✓	PressView technique was particularly helpful for a quick peek into occluded items.
✗	Miniature was found to be distracting and introduced too much visual clutter.
✗	A group accessing technique was missing.
✗	Functional zones was found to be suboptimal for indicating the default awareness level.

Table 3.5: Summary of results from the exploratory experiment

halo is visualized around the occluding physical object as illustrated in figure 3.22 regardless of the number of occluded items. The occluded object remains in its original location (R1) but an icon-sized miniature representation is visualized on the halo. This indicates the rough spatial location of the occluded object relative to the occluding physical object and provides additional information about the occluded object.

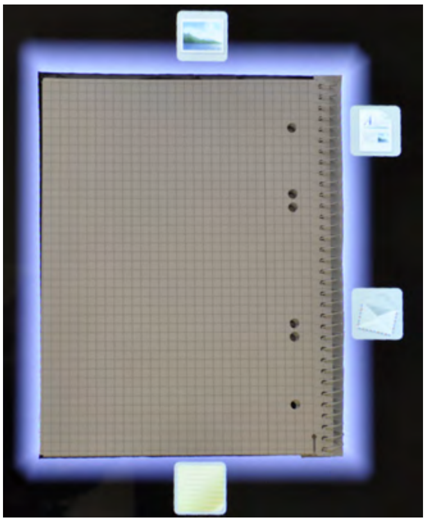

Figure 3.22: Halo visualization and iconized representation of the occluded objects around it

Furthermore, drag-out access with halo visualization is improved in a way that it allows users to resolve occlusion at different scopes (R6). Icon provides individual access to a digital item. In order to access a group of occluded objects users can drag out any part of Halo. In addition to supporting access to occlusion groups and individual occluded objects, we designed and implemented a lightweight way to obtain an overview of all digital items on the tabletop (R1). The *Global Expose* mode is triggered by a five-finger gesture on an empty tabletop area as illustrated in Figure 3.23 (e-f). It automatically moves and scales all occluded objects to become visible. The user can grab any digital object and move it to a new location. In order to avoid permanent clutter, upon releasing the fingers, objects that were not moved snap back to their original, occluded locations.

The halo visualization is inspired by the non-interactive Glow technique introduced by Javed et al. [Javed 2011]. In contrast, both our halo and the individual proxy designs act as interactive widgets to preview and access occluded objects. Icons provide access to individual occluded objects (R1). In order to access all underlying objects users can also drag out and enlarge the halo (cf. figure 3.23c,d). The halo visualization enables global awareness of occluded objects as well as access to occlusion groups and individual objects, while minimizing visual clutter.

3.5.2.2 Remote Accessing

Based on participants' feedback, we improved the remote zooming gesture and adapt it to perform in combination with the halo visualization. While touching a halo or an individual icon with one hand, the other hand performs a pinch gesture on any other area of tabletop display as depicted in figure 3.24. The pinch gesture opens an interactive preview of the object or the occlusion group, offering the same functionality as the direct technique described above. By pinching the fingers apart, the preview gets larger and shows more detailed information. Once the preview is opened, the user can release the finger on the icon to interact with the preview or to reference to another halo. The preview disappears once the pinch gesture is released. Note that with the improved remote access gesture users not only can temporary inspect an occluded object but access it on any empty area on the tabletop. Moreover the refined version enables a remote group access (cf. figure 3.24b), which was not possible with the previous remote zooming gesture.

3.5.2.3 System Prototype, Improved

We improved our system prototype by using the state-of-the-art tabletop device. We reimplemented our prototype, which runs on a Samsung SUR40 with an active

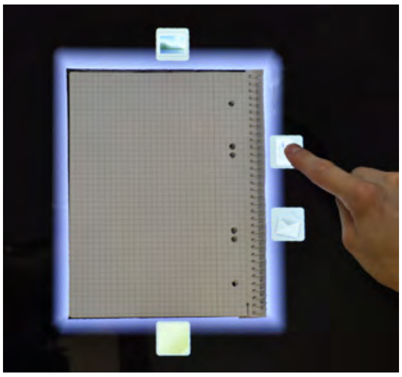

(a)

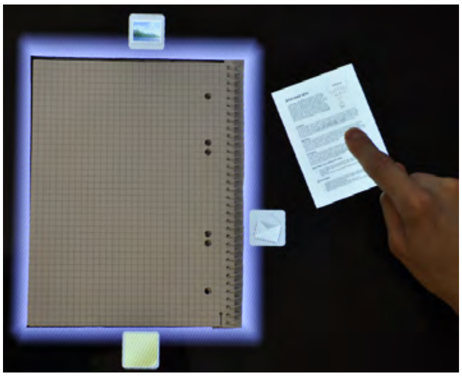

(b)

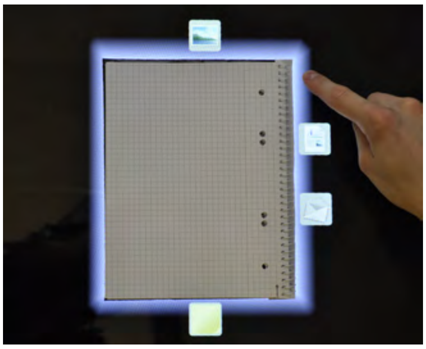

(c)

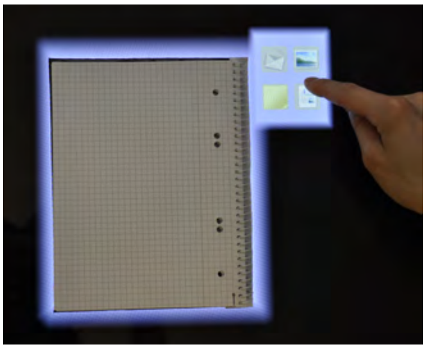

(d)

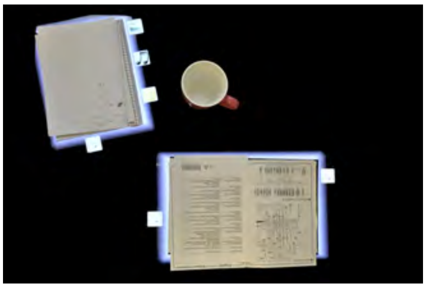

(e)

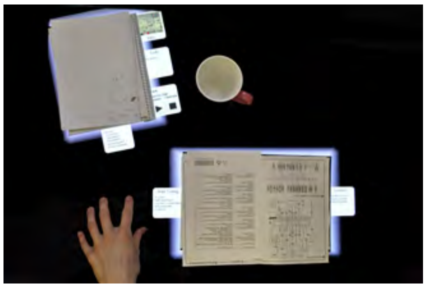

(f)

Figure 3.23: Accessing occluded objects at different scopes: (a and b) individually, (c and d) group access, and (e and f) showing the global expose technique for getting a global overview

(a) (b)

Figure 3.24: Remote accessing technique – users can select either (a) an individual proxy or the halo for an individual proxy or (b) a group access, accordingly.

display area of 40 inches and resolution of 1920 x 1080 pixels. The system also includes a computer with a dual-core processor with 2.9 GHz. In contrast to the previous system which utilizes a camera for touch tracking and marker recognition, it uses the Microsoft PixelSense technology [4] for touch and physical object recognition on its surface without the use of cameras.

The PixelSense technology (cf. figure 3.25b) basically uses a back diffuse illumination approach, which with the aim of a high-resolution grid of integrated sensors can beam as well as observe IR lights reflecting back from the fingers and objects footprints. As it can be seen in figure 3.25, the SUR40 offers a relatively compact (only four inches thick) and more table-like form factor so that users can conveniently sit around it. This feature makes SUR40 particularly suitable for user studies.

The object recognition on the Samsung SUR40 system is realized through fiducial tag markers. Thanks to the Microsoft PixelSense technology, the tag markers can be seen and decoded with the size of only 19 x 19mm. The marker basic coding scheme is composed of four reference points – namely, one large point in the middle of the marker for extracting the position and three guide points to specify the orientation of the tag. In addition, each tag contains from zero to eight data bits that define the tag value. These data bits are white circles (0.075 inch radius) that are clockwise centered around the tag.

We reimplemented the occlusion-aware techniques described above in a fully functional system written under .NET framework in C# programming language. Tag markers enable tracking physical objects on the tabletop surface. Although

[4] http://www.microsoft.com/en-us/pixelsense/pixelsense.aspx

(a) PixelSense Technology

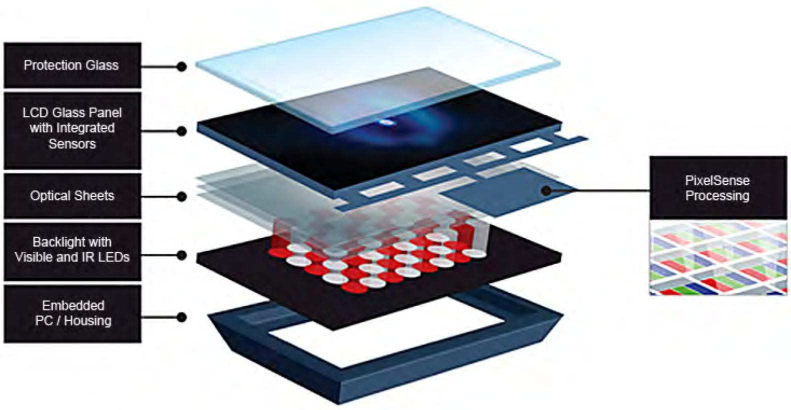

(b) Samsung SUR40

Figure 3.25: Samsung SUR40 tabletop system which uses Microsoft Pixelsense technology

this requires augmenting objects with tags in advance to recognize physical objects placed on the tabletop surface, it aids study of the occlusion aware techniques and avoids any additional top-mounted instrumentation.

The system detects percentage of occlusion. Specifically, when more than two-thirds of a digital object is covered, it is recorded as being occluded. In order to position the icons on the edge of a physical object we implemented a basic yet responsive and real-time approach. For each occluded digital object, it first computes if the left or right side of the physical object is nearer to the digital object, leading to either (left, top, bottom) or (right, top, bottom) as possible placements.

In the next step, the distance to the top edge of the physical object is compared to the one to the left border, if it is smaller, then the object gets visualized at the top border. if not, the same is done for the bottom edge. If neither of them was closer, it is finally visualized on the right. For placement at top or bottom edge, the original x-coordinates are retained and y is set so that the icon is above or below the object; for left and right edge placement, the same is done but with y value retained and x value set accordingly.

The software also recognizes two-handed organizational gestures. Basically, the touch points is grouped into gesture objects based on their proximity when they first appear. The gesture object includes: *number of fingers* in the gesture, *speed*, *duration*, *acceleration*, *angle*, *size*, plus a *state* that can be set to active (i.e., touch points still present) or inactive (i.e., gesture ended, this state is seen once by all subscribed objects and then the gesture object is freed).

Then, a visualizer component employs a simple decision tree to determine the gesture means. For instance for teleport, if the visualizer sees an inactive (terminated) one finger gesture with a move distance of at least x mm and a duration less than y ms, then it is considered a swipe. The visualizer then asks the gesture processor if there is any active three finger gesture on the table; if so, the object gets moved there. Another example is rotating: if the visualizer has an active two finger gesture on it, it adds the rotation delta of this gesture object to its own rotation value. In this fashion, all gestures are detected independently and perform more robustly.

3.5.3 Study2: Controlled Experiment

The study revealed that moving or lifting physical objects is a natural practice that can be performed even faster than dragging out occluded objects using our techniques. Particularly, participants preferred replacing occluders with lightweight and easily movable objects. This motivated us to further examine which strategy is more suitable and performs faster under what circumstances and how it correlates to properties of physical occluders. Therefore, one aim of this user study is to formally compare this well-established strategy with our occlusion-aware techniques for providing awareness and visibility of occluded items. The objective is to compare performance for a set of common tasks in both low- and high-clutter hybrid tabletop conditions.

We selected a conventional tabletop system, in which users can move and lift occluders to cope with physical occlusion problems, to serve as the baseline condition. While it may seem that the comparison with this baseline favors the proposed techniques, in the previous study we observed that moving and lifting everyday objects is easily performed and can be as fast as dragging out occluded objects. We opted not to compare our occlusion-aware concepts with prior systems or techniques (e.g., [Furumi 2012, Javed 2011]), because moving and lifting physical occluders is the most direct and most intuitive strategy to cope with occlusion and this has not been examined in previous work. Thus, it is valuable to formally compare these two techniques and identify their respective advantages and disadvantages.

We also added a third restricted condition, in which participants could not move occluders to simulate interaction with heavy or bulky occluders that cannot be easily moved (e.g., screens, desk lamps, plants, or piles of books, folders, and documents) or where the specific location of the object serves other purposes (e.g., a keyboard positioned in a comfortable location for typing). In summary, we compared three conditions – namely,

Classic conventional tabletop system, wherein users could freely move and lift physical objects within the tabletop screen space, although stacking physical objects onto each other was not allowed. Users could manipulate digital objects using common touch gestures (move + pinch).

ObjecTop Move is the same as the classic condition but occlusion-aware techniques could be used.

ObjecTop Immovable constrained ObjecTop condition with immovable physical objects. Users were not allowed to move or lift physical occluders.

Compared to conventional tabletop systems, we anticipated that both *ObjecTop* conditions would improve the overall performance of finding and structuring digital objects, particularly in high-clutter hybrid settings. Our specific hypotheses are

H1: Finding digital objects using occlusion-aware techniques in both *ObjecTop* interface conditions will be faster than with the classic interface.

H2: Both *ObjecTop* interface conditions will outperform the *classic interface* in organizing digital objects.

H3: The *ObjecTop Move* condition will significantly decrease interaction with physical objects compared to the *classic condition.*

3.5.3.1 Method

Participants

A total of 17 unpaid participants were recruited from our institution. One of them was female and all were right-handed ranging from 23 to 38 years old (M = 28.3, SD = 4.3). All had experience with touch interfaces but only two had experience with interactive tabletops. The study took place in our lab environment.

Study Setup

We used our system prototype described in section 3.5.2.3 consisting of a Samsung SUR40 tabletop systems. It allows participants to comfortably sit around and provide for more ergonomic hybrid tabletop setting.

Materials

In both tasks, for physical items we used a thick book, a thin paper notepad, an iPad (in keyboard stand position) in the low-clutter condition. In the high-clutter condition, an A4 paper tray containing a pile of paper and a coffee mug were added. A schematic view of the study setup is depicted in figure 3.26. The digital items were photos, documents, and Post-it notes.

We manipulated the number of digital distracters as well as physical objects to vary the level of clutter. In earlier interviews, five individuals were asked to rank the complexity level of hybrid tabletop settings that included both digital and physical objects. Based on these interviews we defined clutter to be a function of the number of digital (ND) objects on a tabletop and the extent to which display surface is covered with physical objects (PO). For our experiment, the low-clutter condition consisted of ten digital objects of which five were occluded using three physical objects (PO 40%). The high-clutter condition involved 15 digital objects of which ten were occluded using five physical objects (PO 60%). There were no digital-digital occlusions.

Tasks and Procedure

The experiment consisted of two tasks: a search and a layout task. In the search task, participants were asked to find and count specific types of target virtual objects among a number of distracters. One to three target objects were randomly chosen. The targets were uniformly distributed across all tabletop zones (both storage and working zones 3.12) and were fully occluded in all trials. For this task, we only enabled system features related to awareness and accessing occluded objects.

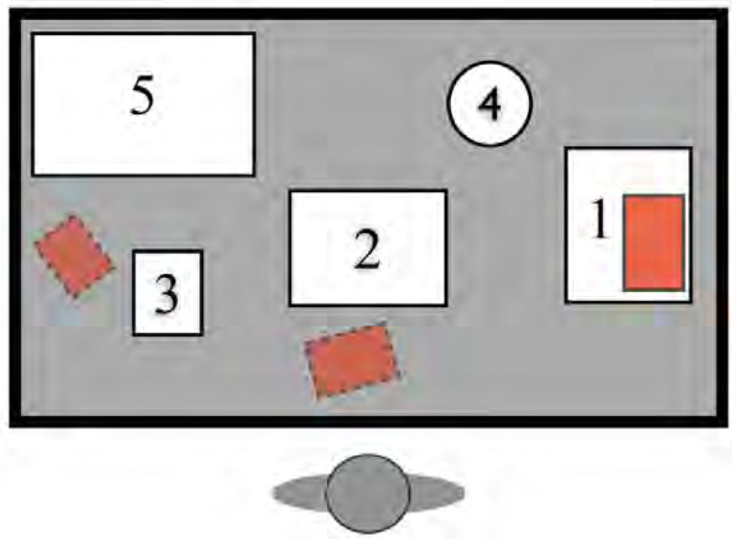

Figure 3.26: A schematic view of the study setup

In the layout task, participants were asked to first find three target objects across a set of physical and digital distracters. They then had to move two of them and place them on two containers displayed with a virtual, faded-out preview of the target objects (cf. figure 3.26 dashed red objects). The third target object had to be hidden under the physical book. To indicate the hiding action, we placed a faded-out printout of the third target on the book during this task (cf. figure 3.26 solid red object). The location of the containers and the physical book were the same in all trials. In addition to occlusion-aware accessing and awareness techniques, we enabled teleport and hiding gestures in *ObjecTop Move* and *Immovable conditions*. For the study, we ignored the hybrid binding and Expose techniques because of the lack of equivalents in the touch user interface.

The experiment employed a: three interface (*Classic, ObjecTop Mov and Immov*) $\times$ 2 task (*Search and Layout*) $\times$ 2 clutter level (*Low and High*) within-subjects design. Our main dependent measurement was trial completion time. It was measured by pressing the spacebar of a keyboard placed on the margin of the tabletop and within arm's reach of participants.

In the search task, participants advanced to the next trial regardless of whether their answer was correct or not. We recorded all actions and movements of digital objects in a time-coded log file for subsequent analysis. The perceived workload of each interface condition was measured with the standard NASA-TLX questionnaire [Hart 2006]. The questions asked in the questionnaire are showed in A.

Each participant completed the three interface conditions. Order of presentation was counterbalanced. In each condition, the search task was performed first, followed by layout task. Each task trial block consisted of two low and high trials repeated five times each. The trial representation within each session was randomized. High or low scene was the same across conditions and participants. The appearance of low- and high-clutter trials was also the same across all interfaces and participants.

Prior to each interface condition, participants received a short introduction about the system and task and were asked to perform it as quickly and accurately as possible. In *ObjecTop Move* and *Immovable* conditions, participants had ten minutes' time to practice the techniques. Before starting each interface condition, there were four training trials. After the completion of each interface, a semi-structured interview was conducted to gather subjective feedback. All sessions were videotaped. Each session lasted approximately one hour.

3.5.3.2 Results

Excluding training trials, there was a total of 1020 trials (17 participants $\times$ 3 interface conditions $\times$ 2 clutter levels $\times$ 2 tasks $\times$ 5 repetitions). We had to repeat 84 trials due to the system's failure to recognize tags.

We processed the data to identify and remove outliers, which are those trials with a task time was greater than three SD away from the mean time. We also identified mistrials which are search task trials that participants answered incorrectly. In total, we removed 18 trials from the layout task and 15 from the search task. In the following analysis, we used Bonferroni corrections for all post-hoc tests. We found that presentation order of the task blocks had no significant effects on the task time (no learning effect). Our primary interests are how the interfaces performed in terms of time performance, number of interactions with physical objects and the perceived task load. Each is presented in turn below.

Task Time Performance

A two-way repeated measures ANOVA on the factors of the tabletop interfaces and

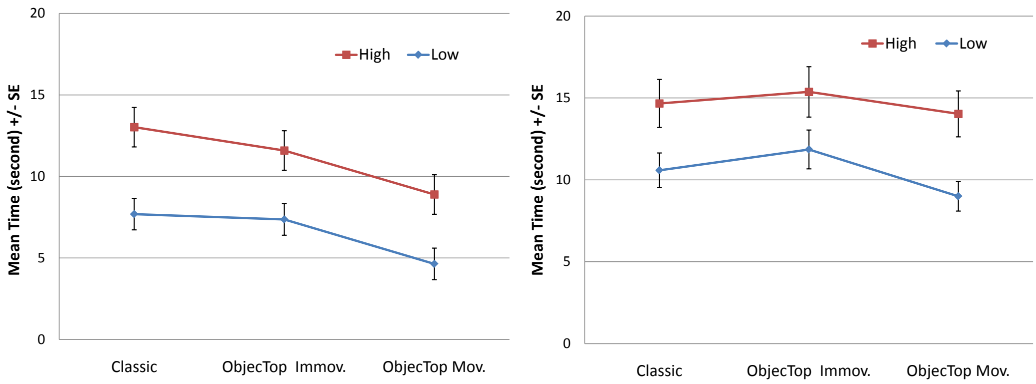

(a) Search task (b) Layout task

Figure 3.27: Task time performance

clutter level for the search task identified main effects for the tabletop interface as well the clutter level (F(2,32) = 2,17, p<.001 and F(1,16) = 4,785, p<.001 respectively). Pair-wise comparison revealed that ObjecTop Move was significantly faster than the classic and ObjecTop Immoveable conditions (p<.001 and p<.01 respectively). We found no interface×clutter interaction (F(2,32)= 3,043, p=.069). Further examination of each clutter level showed that ObjecTop was significantly faster than both classic and ObjecTop Immoveable for the low-clutter level (p<.001 and p<.01). For the high-clutter level, both Move and Immoveable ObjecTop conditions were significantly faster than the classic condition (p<.001, p<.05 - figure 3.27, left).

For the layout task, a two-way repeated measures ANOVA identified the interface (F(2,32)=7,133 p<.01) and clutter (F(1,16)=4,93 p<.001) factors as main effects. Pair-wise comparisons showed that the ObjecTop Move interface was significantly faster than the ObjecTop Immoveable interface in both low- and high-clutter levels (p<.05 , p<.05 , see figure 3.27, right). We found no other significant differences.

The results of trial time completion revealed that search time using the occlusion-aware techniques in addition to moving objects (*ObjeTop Move* condition) was significantly less than with the classic tabletop interface. This supports our first hypothesis. However, in layout task, the time needed to recall and perform the gestures was longer, particularly in ObjecTop Immoveable compared to classic condition. Thus H2 is not supported.

Interactions with Physical Objects

We analyzed the log data and calculated number of interactions (moving or lifting)

task	factor	F-statistics	sig.
search	interface	$F_{(1.16)} = 100.461$	P<.001
	clutter	$F_{(1.16)} = 15.030$	P<.001
	interface x clutter	$F_{(1.16)} = 0.365$	P=.554
layout	interface	$F_{(1.16)} = 20.659$	P<.001
	clutter	$F_{(1.16)} = 26.851$	P<.001
	interface x clutter	$F_{(1.16)} = 4.857$	P=0.133

Figure 3.28: Statistics of physical object interactions

with physical objects for ObjecTop Move and classic interface conditions. In both tasks, we found main effects for both interface and clutter factors. Statistics are summarized in Table 3.28. Post-hoc comparisons for both tasks showed that ObjecTop significantly reduced the number of interactions with physical objects in both low- and high-clutter levels (each with $p<.001$; see charts in figure 3.29).

Since with the ObjecTop Move interface participants could cope with occlusion either by moving or using system support, we analyzed log data to see how individual physical objects were manipulated in this condition. We filtered out movements (less than five mm) of physical objects caused by inappropriate recognition of markers. A summary of the number of interactions with each object is given in Table 3.30. The book and paper pile tray were moved much less frequently than the notepad and coffee mug.

Overall, the results demonstrate that occlusion-aware techniques helped participants to cope with occlusion, requiring fewer physical interactions in both search and layout tasks. Therefore, H3 is supported.

Perceived Task Load

We collected the perceived workload data using a scale of 1-20 (1 is least effort) for various types of workloads: mental effort, physical effort, temporal demand, performance, overall effort and frustration. A one-way repeated measure ANOVA found the interface to be a main effect on physical effort, overall effort, frustration and performance (see figure 3.31). Post-hoc pair-wise comparison with Bonferroni correction revealed that the classic condition caused significantly more physical ($M=14.5$, $SD=3.6$) and temporal ($M=12.64$, $SD=3.06$) effort than the ObjecTop conditions (both with $p<.001$). This resulted in significantly higher overall

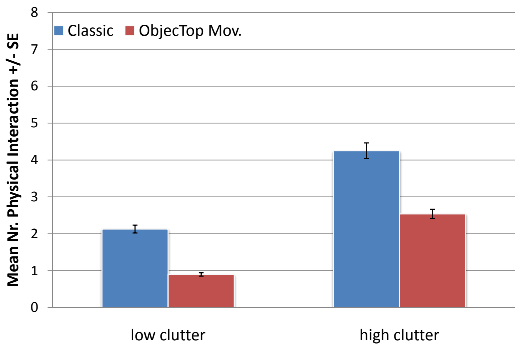

(a) Search task

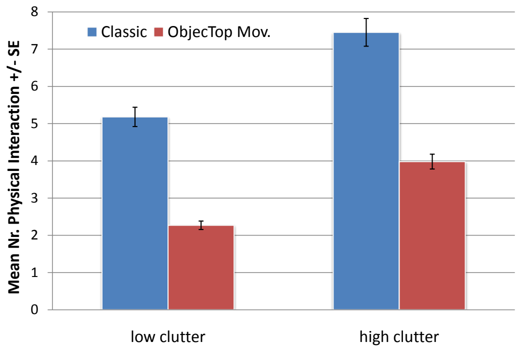

(b) Layout task

Figure 3.29: Interactions with physical objects

id	items	size	low			high		
			Nr.	Mean	SD	Nr.	Mean	SD
1		23 x 19 cm	43	2.8	3.2	34	1.9	2.3
2		24 x 18 cm	53	3.5	3.6	28	1.8	1.6
3		18 x 13 cm	84	5.6	3.4	65	4.3	4.6
4		5 cm (diameter)	-	-	-	48	3.2	2.1
5		30 x 22 cm	-	-	-	35	2.3	2.2

Figure 3.30: Physical object usage in the ObjecTop Move Condition. (Nr.: the number of physical interactions with each object, mean: mean value, and SD: standard deviation of Nr.)

effort (M=13.5, SD=2.4, p<.001) with the classic interface. In addition, the classic tabletop condition resulted in significantly higher reported frustration (M=12.9, SD=2.48, p<.001). We found no other significant differences.

Table 3.6 summarizes the main results of this study. Based on these results we draw several conclusions and discussions presented in the next section.

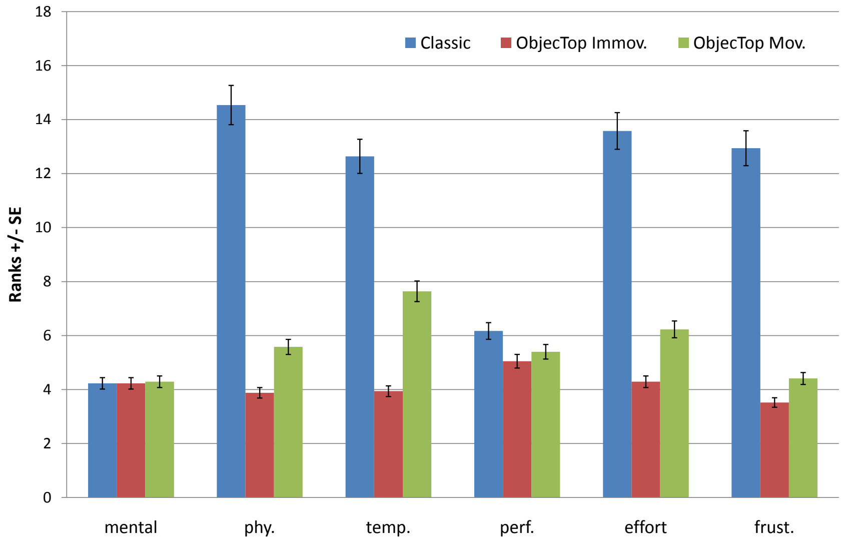

Figure 3.31: Perceived task load of each interface condition

✓	ObjecTop performed significantly faster than the conventional tabletop for searching task.
✓	Physical interaction with occluders (particularly with heavyweight ones) was significantly fewer in ObjecTop condition.
✓	The book and paper pile tray were moved much less frequently than the notepad and coffee mug.
✓	conventional tabletop system caused significantly more effort and frustration than ObjecTop system.
✗	Organizational techniques required time to recall gestures.

Table 3.6: Summary of results from the controlled experiment

3.5.3.3 Discussion

What is the better strategy? Dragging out occluded items or Moving physical Occluders?

The study demonstrates that the physical properties of occluders significantly influence how participants deal with occlusion. In order to cope with occlusion, they prefer to move or lift physical objects that are lightweight, easily graspable and have a small footprint (e.g., the coffee cup or notepad). However, for complex and heavier objects (e.g. the book or pile of paper) they preferred using occlusion-aware digital access.

In addition, physical objects that are placed in the storage areas (outer areas on the tabletop surface at a distance from the user) are less frequently moved, and participants mainly used our techniques functionality to access occluded objects, even though they had the option to lift and move occluders. Our techniques enabled participants to employ an efficient combination of strategies based on the physical properties of the occluder and current layout. This resulted in less physically demanding interaction with objects, lower perceived effort and frustration while manipulating digital objects.

Participants mentioned that the system provides a sort of “structured access” (P11) to the occluded objects and “decreased number of unnecessary interactions with physical and digital objects” (P1). P10 said, “All physical edges of occluders become a kind of taskbar where I can spatially organize and minimize my digital objects.” We observed though that once a large number of digital objects were occluded under a physical one, participants preferred physical access by lifting or moving the occluder regardless of its size. A possible reason for this might be that dragging out a large group of digital occluded objects in our system design required more free space around the physical occluder.

The study findings indicated that there are conditions in which each of the resolving strategies (physical interaction and digital pulling out) is preferred and practical. In addition, we suggest a third strategy to resolve occlusion. It employs manipulating physical objects and digital access concurrently. As a concrete example, imagine that as soon as the user partially lifts the occluder to peek under it, occluded digital objects slide out visualized as icons. In this way, space created by partial lifting of the occluder can be used for representing occluded digital objects. This helps with the lack of adjacent space problem discussed above.

Quality of the system reaction

In interviews, we asked participants about the amount of visual clutter introduced by halos and icons and the extent to which they found it distracting. Almost all participants (15 out of 17) agreed that the amount of added visual clutter is acceptable in a low-clutter level. P16 said, “It helps somehow to make my digital information space even more tidy by quickly hiding, moving and iconizing the objects. It results in [a] more tidy setup.” However, in a high-clutter level, participants were skeptical about the amount of visual clutter and proposed some improvements. Five participants commented that the system reaction to and level of awareness about the occluded objects should be somehow under the user’s control and manually adjustable (e.g., for the whole tabletop or even one occlusion group).

Where to visualize occluders on the halo?

The interactive halo design places an icon on the halo closest to the center of the occluder's footprint. The assumption underlying this design decision was that it implicitly conveys spatial information about the location of the occluded object. However, during the study we observed that participants mainly used the group access gesture (N=443, mean=29.53, SD=17.3 per participant in the whole study) to check all objects occluded underneath one occluder, instead of looking at just one occluded object.

In interviews, 13 participants stated that the icon placement on the halo should be toward the user. They emphasized that this is particularly important for occlusion groups that are located on distant areas of the tabletops or for volumetric objects that occlude more tabletop space from the user's view than just their footprint (e.g., a laptop while the lid is open). Doing this may decrease the effort required to look for icons. However, several participants commented that they prefer our original design for smaller and lightweight occluders, since a representation close to original location of the occluded items provides a hint about its approximate position and makes it easier to move the occluder in a way to best access the occluded object.

Gradual and remote access

With respect to gradual access, we observed that participants rarely used intermediate detail levels while accessing individual items. Analysis of the videos revealed that accessing the various detail levels of an individual item requires a fine-grained action but was mapped to the course dragging-out gesture. Applying hysteresis to maintain each detail level for a longer duration while dragging out may alleviate this problem.

Participants, however, frequently used zooming for the group access since it enabled a peek under the space covered by physical occluders. Remote access (remote zoom) was mainly used in the ObjecTop Immovable condition, since participants were not allowed to move physical objects. Using remote gesture, participants occasionally resolved occlusion groups that were located on more distant tabletop areas (storage areas).

Nevertheless, in interviews participants generally appreciated remote access and commented that it can be useful when the arrangement of the physical objects should be maintained (for instance, a spatially sorted arrangement of a set of documents). This can be expected to be more frequent in everyday information organization tasks.

Organizing the hybrid workplace

We analyzed how participants interacted with digital objects in the layout task. We observed similar patterns in the ObjecTop Move and Classic conditions: participants organize digital objects using conventional touch gestures and manually freed the path while dragging the digital item to its target position. In contrast, in ObjecTop with immovable occluders, participants used either the gestures (teleport and hybrid hiding techniques) or moved digital objects around the occluders. This caused longer task completion time in the ObjecTop immovable condition. Discussion with participants revealed that the organizational gestures (Teleport, hybrid hiding, and binding) are only practical for interaction with digital objects located on storage areas of the tabletop.

We observed that participants used the Teleport gesture more (N=141, mean=9.4, SD= 5.3) than the hiding technique (N=36, mean=2.4, SD= 2.5). A common pattern was that they first moved digital objects near an occluder using Teleport and then simply hid them by pushing them beneath the occluder. Discussion with participants revealed that hybrid hiding and binding gestures needed more visual focus and hand coordination. In contrast, participants mentioned that the teleport gesture is much easier to perform since "you just need to find or make an empty spot to place the tripod gesture and then that's it" (P6). Five participants even suggested extending the teleport gesture to be performed on physical occluders for hiding objects underneath them by placing a tripod-like, three-finger gesture atop the physical object.

Occlusion-aware techniques in multi-user settings

Although presented techniques are implemented for and studied in a single-user tabletop setting, they can be easily extended to support collaborative scenarios (e.g., meeting scenarios where multiple users sit around a digital tabletop working with multiple physical and digital objects). Once occlusion happens in the shared space of the tabletop surface, the halo and icon representation can be extended in a way that provides orientation-invariant access to the digital objects. As a concrete example, the halo can provide strong visual clues on each corner of the occluder to enable group access to occluded objects from four directions. Moreover, since the system knows the position of the user dragging out the halo, it can visualize the various detail levels of occluded objects in an appropriate orientation.

The occlusion-aware organizational gestures can also be extended to cooperative gesturing [Morris 2006] in multi-user settings. This may enhance collaboration as well as facilitate reachability and access across large surfaces. Future work is needed to investigate issues related to cooperative gesturing in hybrid tabletop settings.

3.6 Conclusion

In this chapter, we addressed the integration of everyday physical objects with interactive tabletops. We particularly focused on the challenges stemming from the placement of physical objects on the tabletop surface giving rise to partially or fully covering digital items. We argued that physical occlusion on tabletop pose challenges for both display and interaction with digital items. Users lose awareness of occluded digital item and cannot effectively access them. Since it is reasonable to expect that tabletop display surfaces will increasingly become part of everyday life, understanding how to provide practical solutions to the challenges of occlusion is important.

To this end, in this chapter we contributed an occlusion management framework called *ObjecTop* for hybrid tabletops, which consists of a number of occlusion-aware concepts supporting awareness, access, and organization of digital items. We grounded the design of ObjecTop in two ways: first, by analyzing results of the hybrid media study presented in the previous chapter (2.2) through a physical occlusion perspective and second, by conducting an exploratory study in which we examined not only occlusion created by printed documents but also other ordinary objects. We then compiled all findings into a set of requirements to be considered while designing ObjecTop concepts and techniques.

To assist accessing occluded items, we employ *proxies* as a semantically zoomed-in representation of actual objects. The proxy has two important roles in ObjecTop: as a means of awareness and as an access tool. Based on functional zones on tabletops, proxies provide various levels of detail about the occluded object. Users can access occluded objects by dragging out respective proxies. To support a hybrid strategy for resolving occlusion we also designed a pressure-based accessing technique called PressView so that users can slightly press on the physical object in order to squeeze out the underlying proxies. Furthermore, ObjecTop implements a set of interaction techniques for efficient organization of digital objects across physical occluders. This included techniques for moving, zooming, piling, binding digital objects.

To validate the design of ObjecTop we conducted two user studies. In the first one, we exploratively examined the main features of ObjecTop (i.e., the use of proxies, drag-out and pressure-based techniques) in a realistic hybrid tabletop setting.

Among the other findings, the proxy concept was very well received by users. Regarding different levels of details (i.e., halo, icon, miniature) offered by proxies, users found Icon provides reasonable tradeoff between information that it conveys about the occluded objects and the visual clutter that it produces. Users also found PressView as a lightweight means for quickly looking at occluded items. Moreover it was found that users frequently moved or lifted physical objects to resolve occlusion. This behavior particularly occurred once occlusion was created by lightweight and easily movable objects – such as a coffee cup or printed documents.

The results of this study motivated us to drill down through key aspects of our system in a second study in which we formally compared ObjecTop with the natural way of coping with occlusion (i.e., moving and lifting physical occluders). The second study systematically confirmed that ObjecTop is faster for finding objects, decreases interaction with physical objects when resolving problems of occlusion, requires less effort, and is less frustrating.

ObjecTop interaction facilities are examples of providing what one might term an informational physics for digital elements [Bederson 1996]. Just as the simple physics of snap-dragging is valuable when positioning items on a grid, one can envision a variety of informational physics designed to support specific tasks while being sensitive to the current activity context and the hybrid collection of digital and physical objects. The challenge is to support information-based activities in ways that increase not only efficiency but also enjoyment of interaction.

Chapter 4

Exploring Physical Interactions for Paper-like Displays

Contents

In the last two chapters, we approached the integration of knowledge work physical settings with interactive tabletops. We investigated how ordinary objects in such settings can be effectively combined with virtual objects displayed on the tabletop surface thus, supporting knowledge workers in so-called hybrid tabletop settings.

In chapter 2, we showed how to design interfaces for hybrid physical-digital document piling. We exploit the flexible arrangement of multiple printed documents for assisting manipulation of digital documents in a hybrid pile. Chapter 3 then focused on physical interaction concepts that are influenced by the physical layout and arrangement of ordinary objects on the tabletop surface. Physical occlusion (i.e., physical objects hide partially or fully digital ones), turned out to be one of the most striking challenges that needed to be addressed. In response to that, we contributed an occlusion management framework to mitigate challenges of physical occlusion. In general, both chapters investigated physical interaction concepts to support knowledge workers from the *table-centric* perspective.

In the present chapter, we aim to support knowledge work practices from the *paper-centric* perspective. As it is mentioned in the introduction of this thesis, *paper* is an indispensable component of knowledge work activities due to its unique advantages: it is flexible, thin and lightweight so that can provide variable screen real estate that suits the current context of use. With the advent of thin-film display technology that potentially provide ultra thin and deformable displays, many of such physical properties that until now were unique to paper can be incorporated for interaction with digital content. Multitouch displays might ultimately become so thin that they can be arbitrarily folded and rolled while featuring high-resolution display both on the front and the reverse sides. This does not only offer a high level of portability but also opens up novel interaction possibilities that leverages physical resizing of display as an input modality.

In this chapter of the thesis, we investigate how future hand-held computing devices featuring ultra-thin displays can support knowledge worker practices. Here we particularly leverage folding and rolling as navigational and space saving techniques for designing physical interaction concepts that support more embodied, expressive, and joyful interactions with digital contents. In contrast to other paper-like manipulation techniques – such as bending – folding and rolling are particularly of interest because they modulate size and thus form factor of a flexible display. Based on an extensive literature overview, we set out a number of design rationale for future *resizable displays*. We then propose two novel device concepts – namely, *FoldMe* and *Xpaaand* – in which we explore the design space of foldable dual-sided and rollable displays, respectively. We design interface concepts that leverage physical resizing of displays through folding and rolling for accomplishing a set basic tasks with digital contents.

In summary, contributions of this chapter are

- two novel device concepts that feature resizable displays along with their physical interaction design spaces
- novel interface concepts and interaction techniques
- simulation environment that emulate the look and feel of paper-like displays
- evaluation of concepts through two exploratory user studies using functional prototypes

This chapter is organized as follows. We begin by providing background and reviewing the previous work in relevant research fields reported in 4.1. Drawing upon literature analysis and our design goals, in section 4.2 we depict two novel device concepts that feature flexible displays for dynamic resizing of screen real estate. Section 4.3 proposes a set of user interfaces for such devices along with several concrete interaction techniques. After describing our implementation details and device prototypes in 4.4 we report on two user studies documented in section 2.5 in which both device concepts and their user interfaces are evaluated.

> **Contribution Statement:** Most of the work presented in this chapter is based on and has been published in [Khalilbeigi 2012a, Khalilbeigi 2011]. I am the first author on these publications. I have initiated and lead both FoldMe and Xpaaand projects. My co-authors have also contributed significantly. Master students, Wolfgang Kleine and Jan Riemann, have built and implemented many aspects of the simulation environment as well as applications running on the devices. Roman Lissermann and my supervisors, Jürgen Steimle and Max Mühlhäuser, have contributed to the design of the interaction concepts and to writing the papers.

4.1 Background and Related Work

As mentioned in the introduction chapter, despite the prediction of paperless office, paper still exists and its usage is ubiquitous as it has been documented by Sellen and Harper [Sellen 2003]. They identified a number of key affordances that explain the infinite battery life of paper documents. As opposed to section 2.1 where we analyze the well-established practices of paper – such as flexible piling and rearranging –, here based on the literature [Sellen 2003, Sellen 1997], we systematize the popularity of paper documents along the following four characteristics:

1. Physical Flexibility: paper is flexible, low in weight, and thin; thus it affords a high degree of ecological flexibility. It can be easily manipulated by hand while it is rugged. Moreover, due to such physical specifics, it allows users to easily carry it around and, thus, is highly portable.

2. Malleable Shape and Size: paper is flexible in shape and size. This means that one can easily adjust shape or size of paper documents so that it suits the context of use. For example, newspaper can be folded to ease mobility and where space allows, it can get unfolded to facilitate comfortable reading. Therefore, paper provides more ergonomic usage and allows for malleable screen real estate that can be adjusted based on the current context of use.

3. Tactile and Kinesthetic Feedback: through its tangibility paper offers rich tactile-kinesthetic feedback. People use both hands to interact with physical documents for searching, skimming and interleaving navigation with other activities. Moreover, through the thickness people can approximate the length of document and the remaining number of pages.

4. Multi-display Experience: paper is cheap and has many physical pages that offers a rich multi-display experience. This characteristics facilitates local as well as global document navigation while for instance active reading [Marshall 2005b, Chen 2008]. Local navigation is when users consult materials on and around the current page – for example, by flipping or partially folding the page to glimpse on the previous or next pages. Global navigation spans across the whole document – for example, leafing through book to find a specific point. In case of multiple paper documents, users can flexibly arrange them in the physical space for ease of cross referencing, sorting, prioritizing or expressing their internal relationship.

Beside above and the other advantages, paper has one major downside: the information on paper is fixed to its medium (*fixedness of information* [Sellen 2003]). This means that information on paper cannot dynamically be rewritten, revised or reformatted. Furthermore, paper must be used locally and its information cannot be accessed remotely. The emergence of flexible display technologies – such as OLED (organic light-emitting diode) or electrophoretic (E Ink) [Co 2008] – potentially aims to overcome these limitations by bringing digital computation to displays that closely approximate the look and feel of traditional paper. In the near future, displays or ultimately portable devices featuring such flexible displays can be folded similar to the traditional paper for navigation or they can become so thin that people can roll them similar to the traditional paper rolls. In this way, dynamic information can be displayed and manipulated on *paper-like displays* allowing for:

- more ergonomic usage and high level of portability.
- malleable devices that adapt their device shape to the functionality
- paper-like interactions that draw upon the affordances of paper supporting more expressive and varied ways of interactions with and representation of the digital information.

While the state-of-the-art technology of flexible displays is still under development [Crawford 2005], there is a growing trend in studying user interfaces and interaction techniques for flexible computers that are based on the deformation of displays. Here, two streams of research can be distinguished. First, those prior studies that investigate device deformation as input separated from output (display). We believe that the research in this field can extensively inspire the design of interfaces for flexible computers due to the similar style of interaction that is based on the deformation. A second stream of research focuses on user interfaces for flexible displays where input and output are on the same device. In the following, we review related work in these two areas. Furthermore, our foldable device concept (FoldMe) proposed in this chapter draws upon the previous work in the area of multi-slate device form factors, which is discussed at the end.

4.1.1 Deformable User Interfaces

In early studies, researchers investigated various possibilities of deformable interfaces with simple technology. Fishkin and Harrison et al. investigated manipulative user interfaces in which physical manipulations are directly integrated with the device or object – such as a small PDA – that is being controlled [Fishkin 2000, Harrison 1998]. Inspired by the real world, they mimicked the same style of physical interaction to manipulate digital information represented on devices.They simulated the turning of pages by pressing sensors at the upper corners of the frame of a device. They also explored the use of tilting to scroll menus, using the metaphor of turning cards in a binder. ShapeTape [Balakrishnan 1999] probably was one of the first deformable input device that facilitated direct manipulation of curves and surfaces in a computer. The device is basically a 48 x 1 x 0.1 cm rubber tape that senses bending and twisting using two fiber optic bend sensors. Users can control a 3D virtual curve or surface shown on a computer by deforming the input device.

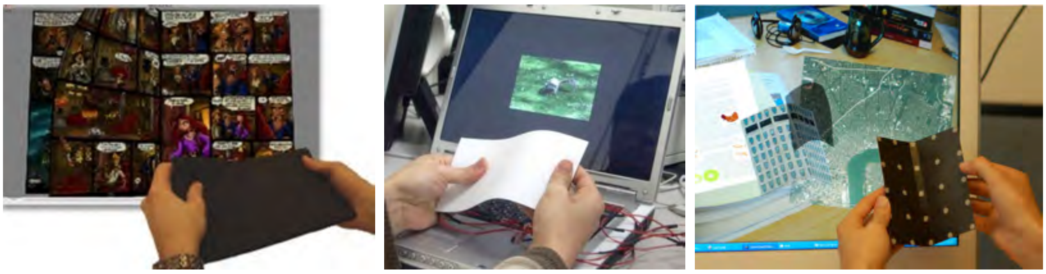

(a) Twend (b) Bookisheet (c) FID

Figure 4.1: Examples of previous work on deformation-based input for manipulation of digital content displayed on separated displays

With the advent of advanced sensor technologies, researchers started investigating deformation-based interactions with more paper-like input devices. Twend [Herkenrath 2008] is a hardware prototype that can be bent and twisted to perform navigational tasks – such as scrolling through the pages of an ebook or a map application. It basically made of a foam layer sandwiched between two thin plastic layers. The prototype measures the bending state using eight optical bend sensors connected to a nearby PC. The prototype (cf. figure 4.1a) allows for bending four edges as well as the whole device along horizontal and vertical axes back and forth. Additionally, it can be twisted into a horizontal wave form with the valley either on the left and right side of the device.

Similarly, Bookisheet [Watanabe 2008] is an interface based on the metaphor of turning the pages in a book. It mimics the natural behavior of leafing through pages by bending a thin plastic sheet in different directions. Using the prototype shown in figure 4.1b users can scroll a photo collection by bending the device or perform finger bookmarking by pressing a micro switch placed near the forefinger position. A comparative evaluation of Bookisheet with other conventional devices (namely, books, arrow keys, and the mouse wheel) showed that it has the same level of performance for retrieval and paging tasks as the conventional devices.

Gallant et al. [Gallant 2008] designed a foldable user interface, which is basically a paper-based foldable input device to control 3D GUIs. The input device shown in figure 4.1c is made of a $5 * 6inch$ cardstock paper augmented with a number of IT retro-reflective markers, the shape and position of which is tracked through an IR camera. The foldable device recognize a plenty of different paper-like deformation gestures – such as leafing, folding, squeezing, or shaking – in order carry out different tasks in the foldable user interface. They can be used to navigate the desktop and to select, stack, sort, annotate, and browse documents.

There also exists a number of studies that investigated how people can leverage physical deformation to interact with digital contents regardless of technological barriers. Lee et al. [Lee 2010] aimed at understanding how users manipulate deformable displays as input devices using artificial deformable displays. They conducted a gesture elicitation study in which they asked participants to devise appropriate gestures for a set of 11 basic commands while using an A4-sized paper, an elastic cloth, and a plastic as device mock-ups. Their results indicated that users tend to apply ordinary actions to define gestures as a metaphor. This resulted in more intuitive gesture with high level of preference among users. Moreover, they found that users preferred consistent gestures in terms of polarity (i.e., when two gestures of the two opposite-but-closely-related commands were paired).

Similar in nature, [Lee 2012] investigated the effects of size on deformable user interfaces by observing people interacting with display mock-ups. They conducted a study in which users were asked to elicit gestures for a set of standard computing tasks using one large (A4) and one small (iPhone) sized flexible plastic. Their results showed that the small-sized device was mainly preferred and users could define create more deformation based gestures with the iPhone sized displays. The study also showed that most gestures were performed horizontally.

Beyond mock-ups, Kinetic Device [Kildal 2012a] is a prototype similar in shape and dimensions to a smart-phone that senses basic bending and twisting gestures without any visual output. Using this prototype, Kildal et al. [Kildal 2012b] investigated how the stiffness and required deformation degree of mobile phone-sized deformable device influenced how precisely users were able to perform force-matching/targeting task. They found out that while device stiffness did not significantly affect task performance, user comfort and preferences were strongly in favor of softer materials. Their results also suggested that nonvisual interaction with deformable devices is viable. Recent study by Ou et al. [Ou 2013] proposed layer jamming as an enabling technology to realize deformable thin sheet interfaces (such as FlexPad [Steimle 2013]) with tunable stiffness. Such sheet interfaces can yield dynamic haptic feedback and shape deformation capabilities.

In summary, the findings of prior studies in this area showed that deformation-based input

- is a promising way of interaction with digital contents
- allows users to obtain rich haptic feedback from the device by feeling its *tangibility* and *configuration* as they interact

(a) Gummi (b) PaperPhone (c) Snaplet

Figure 4.2: Examples of handheld mobile computers featuring flexible display that can be operated using bend gestures

- results in more expressiveness of interaction with digital contents

Inspired by the real world, previous studies have shown that paper-like deformation of devices is a promising input modality for supporting navigational tasks. It was also found that metaphoric gestures adopted from daily behaviors could play a critical role in the deformation-based interactions [Lee 2010]. These findings influenced the design of our devices and interfaces.

4.1.2 Flexible Display User Interfaces

In the last decade, deformation-based interaction with handheld flexible displays drew the attention of many researchers. Gummi [Schwesig 2004] was one of the first devices that leveraged physical deformation of display for interaction with digital contents. The device prototype consisted of a rigid TFT display attached to the center of a flexible Plexiglas base with integrated bend sensors. In this way, the device allows for bending up and downward to accomplish a set of interface operations – such as zoom in and out in the map application, navigating back and forth in websites, or select and deselect UI components. Initial user feedback showed that users were able to quickly pick up the interaction style of Gummi and the interaction techniques could be effectively and enjoyably used to perform a wide range of simple tasks. Similarly, Tajika et al. [Tajika 2008] proposed an intuitive interface for flexible e-book readers that supports turning, flipping and leafing through pages. Similar to Gummi, their prototype also consisted of an inflexible LCD display mounted on a flexible plastic substrate.

PaperPhone [Lahey 2011] showed the first mobile phone featuring a built-in flexible E Ink display that can be controlled by simply bending the corners or sides of the device. A gesture elicitation study with PaperPhone (cf. figure 4.2-b) showed that in general users preferred bend gestures that were conceptually simpler and less physically demanding. The study also identified a strong consensus on the *directionality* of the bend gestures i.e., the coherence relationship between the actual movement of a deforming device in space and the direction in which the information on the display moves. Furthermore, the study emphasized the importance of *orthogonality* (i.e., associating a unique bend gesture to one single action) as well as *consistency* (i.e., using the same or similar gesture to trigger the same or similar command across applications), which guided our interface design.

Using similar hardware setups, researchers deployed flexible displays in wearable computing [Tarun 2011] or desktop setting [Girouard 2012] computing. Tarun et al. introduced Snaplet [Tarun 2011]: a wearable flexible E-Ink display augmented with sensors that allow the shape of the display to be detected (figure 4.2c). When in a convex shape on the wrist, it functions as a watch and media player or when it held flat in the hand it is a PDA with notepad functionality. Users can interact with and manipulate digital content using bend, touch and stylus based input mechanisms. DisplayStack [Girouard 2012] leverages various physical configurations of multiple flexible displays to provide tools for organizing digital documents using piles and stacks of physical display windows.

While the use of actual flexible displays in the aforementioned studies provided for rich user experience, the current display technology supports only slight bending of the device, are quite rigid and tethered with cables. Therefore, some researchers have studied various forms and more detailed deformation of hand-held displays by augmenting the appearance of passive flexible surfaces with image projection. Lee et al. [Lee 2008] demonstrated four different types of foldable and resizable displays using a low-cost tracking system and a projector. They proposed that devices that can be resized and reshaped offer advantages for mobile contexts. PaperWindows [Holman 2005] augments passive paper sheets with an in-place projection of digital contents. Inspired by the natural practices of working with paper, the authors introduced a set of interaction techniques for individual and multiple displays. FlexPad [Steimle 2013] investigated more detailed and fine-grained device deformation using a projection-based flexible display. Using depth images, FlexPad models and tracks the deformed surface in real time. Their findings established that users were able to perform high detail deformation with both shape-retaining and flexible materials quickly and easily. they also found out that flexible material is better suited for application to controlled using simple deformations (e.g., bend whole device or its edges up or down) whereas shape-retaining materials is well suited for use in applications with complex deformations (e.g., asymmetrical twisting).

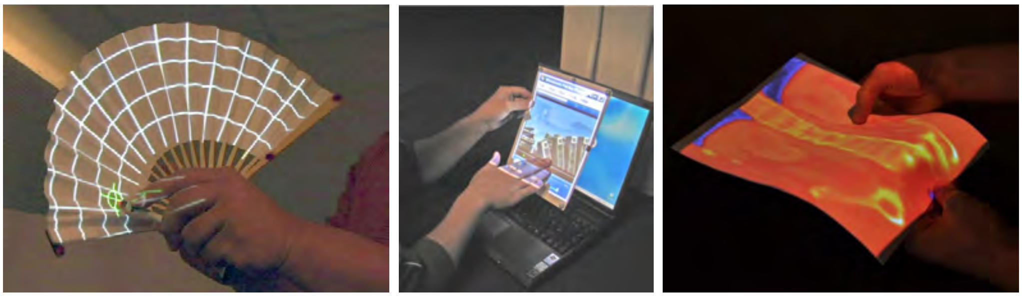

(a) Interactive foldable displays (b) PaperWindows (c) FlexPad

Figure 4.3: Examples of previous work emulating flexible displays by augmenting passive flexible surfaces with top projection

While touch is the dominant form of input for today's handheld computing devices, it is likely that future flexible displays will be multitouch-enabled. Therefore, researchers have investigated user-related issues pertaining to the integration of deformation-based gestures with direct touch input. Dijkstra et al. [Dijkstra 2011] studied the relationship between flexibility of the display and the performance of pointing at specific locations on the display through two user studies. Based on common holds, they identified three structural zones (grip, rigid, flexible) on the surface of flexible displays for which pointing and dragging tasks' performance was formally compared in the second study. The results of the second study indicated that the force characteristics of folds in a flexible display matter. Performance in pointing and dragging tasks is, on average, 12% better in rigid than in flexible parts of the display.

Burstyn et al. [Burstyn 2013] augmented touch input with two types of display deformations – namely, bimanual leafing (bending the side of the display) and one-handed squeezing for depth navigation on flexible displays. Through two user studies they investigated the performance of touch input compared to deformation-based interactions for typical navigational tasks. The results revealed that bend interaction (leafing) is comparable to touch input for navigation through depth-arranged content. However, augmenting touch with the squeeze interaction in parallel, significantly improved targeting performance in a pan-zoom task. Recently, a study by Kildal et al. [Kildal 2013] explored the combination of touch (either on the front or back side of the display) and deformation-based input (bend and twist). Their findings revealed that deformation gestures and touch can be combined successfully as input techniques and the hedonic qualities offered by deformation based interactions are transferred to hybrid interaction techniques that combine deformation and touch.

In summary, the findings discussed above signified that the flexible displays allows for an intuitive and easily understandable way of manipulating digital contents. This is mainly because of the fact that they allows for interaction styles that resemble those used with paper documents. Considerable research in this vein has focused on bending gesture – i.e., curving over the flat surface of a flexible display. Particularly inspired by the study of Lahey et al. [Lahey 2011], it was found that users preferred deformation gestures (concretely bending gestures in the previous work) that are less physically demanding and conceptually simpler. They also suggested to consider three criteria while designing interfaces for flexible displays:

- **Directionality:** the coherence relationship between the actual movement of a deforming device in space and the direction in which the information on the display moves.
- **Orthogonality:** associating a unique bend gesture to one single action.
- **Consistency:** using the same or similar gesture to trigger the same or similar command across applications.

Moreover, prior work that used the passive display approach has shown a practical way of emulating flexible displays. This enables realization of and studying more malleable and paper-like form factors regardless of display-related technological barriers.

While our design is guided by these findings, we extend the previous work by investigating more paper-like deformation-based gestures – namely folding or rolling – which are less explored. We believe that in order to pave the way for effective user interfaces of future flexible displays, each of these modalities need to be systematically investigated. Furthermore, our analysis revealed that such hybrid touch+deformation based input on flexible displays can significantly improve the task performance. Motivated by the prior work we investigate how to effectively combine touch with otherwise bend and twisting deformation-based modalities – namely, folding and rolling.

4.1.3 Slate Display Form Factors

Folding is one of the most common deformation-based gestures used for manipulating paper-based artifacts. While it has not been systematically studied in the context of interfaces for flexible displays, researchers have investigated it as a means to support interaction with rigid multi-slate device form factors. This stream of research is tightly related to and motivated the design of our foldable device concept presented later.

Chen et al. [Chen 2008] introduced an e-book reader that features two displays mounted on two separate slates that are connected by a hinge. The device allows back-to-back and side-by-side configurations of the displays as well as detaching the slates while reading. A set of embodied interaction techniques based on folding, flipping, and fanning of the displays supports both local (i.e., consulting material on or around the current page) and global (i.e., obtaining an overview of a document) navigation in e-books. Their evaluation showed that such dual-display ebook devices have the potential to improve the reading experience by supporting several local navigation tasks better than a single display device. However, embodied interaction techniques turned out to be not useful. Hinckley et al. [Hinckley 2003] proposed to use tablet PCs to span digital content across multiple displays. For example, connection of two tablets results in collating two view ports of displays or triggering full screen preview of an image. Building upon that, they introduced Codex [Hinckley 2009], a tablet computer that features two detachable displays connected via a predefined hinge. Embedded sensors can measure the angle between both displays and the orientation of the device. This provides for a richer design space of different device postures that afford individual, ambient and collaborative use scenarios. While these devices introduced here provide two adjacent displays, each of these displays is single-sided: it cannot display information on its reverse side.

In summary, the previous multi-slate systems showed that dual display device configurations can facilitate local as well global navigation while active reading. Physical interactions – such as page turning and fanning [Chen 2008] – showed their potential to support both types of navigation. Moreover, it was found that ergonomic aspects of the device – like weight and size – play an important role for the success of physical interactions. We extend these work by proposing a device concept that can display information both on the front and their reverse side. This enables an even larger set of device configuration and folding interactions compared to the previous systems. Moreover, we do not only consider single folding gestures – such as page turning and fanning – but also various fold forms, numbers and portions.

4.1.4 Summary

In this section, we reviewed the related work in three different areas pertaining to our main contributions presented in this chapter. Our analysis revealed that deformable user interfaces, be it as an external input device or tightly coupled with output in form of flexible displays, offer a *physical* embodiment of and interaction with digital contents. It was shown that due to the high degree of tangibility in DUI, users obtain a continuous haptic feedback about the state of the interaction. Previous systems and studies showed that DUI allows for interaction styles that closely resemble those used with paper documents. Based on the analysis of previous work, we derived a number of design goals (DG) that that we want to achieve with our concepts and user interfaces. These are listed in table 4.1. The first four design goals (DG1 - DG4) are related to the device concepts presented in the next section. The rest (DG5 - DG7) are criteria that we followed while designing interaction concepts presented in section 4.3.

4.2 Device Concepts: folding and rolling

As mentioned before, the ever growing trend in creating thin-film display technology promises to provide thin and even deformable displays that incorporate many of the physical properties that until now were unique to paper. They might ultimately become so thin that they can be arbitrarily *folded* and *rolled* enabling dynamic modulation of shape and size of the screen of future handheld devices. Such devices potentially can provide a solution to the traditional dilemma of today's fixed-size displays: increased screen size vs. portability. This means that instead of navigating through content spanning outside of display area, users can *physically resize* the display screen real estate of foldable or rollable displays to collate their viewport for seeing more content (cf. figure 4.4). In this fashion, users can dynamically reshape and resize displays to suit desired usage, offer high level of portability.

In contrast to bending, such physical deformation gestures require both hands to grip the display sides or handles on both sides of the device. Thus, based on the Fitts' law model [Fitts 1992], physical resizing becomes inefficient since users have to take off their hands each time they want to manipulate digital content using, for instance, direct-touch input on the display. In our design, we address this disadvantage by exploiting the physical resizing action as means for interaction with digital contents. This opens up novel physical input possibilities that have not been addressed in the previous work.

Design goal		Supported by previous work?	Contribution of this chapter
DG1	Support for physical embodiment of digital contents	◐	Our device concepts feature paper-like flexible displays.
DG2	Support for physical interaction with digital contents	◐	Users can physically reshape and resize the display to interact with digital contents.
DG3	Support for more paper-like deformation-based interactions	○	FoldMe and Xpaaand devices allows for modulation of screen through folding and rolling techniques, accordingly.
DG4	Support for hybrid touch plus deformation input techniques	○	We design different variations of combining touch with rolling and folding techniques.
DG5	Reflect directionality in pairing gestures with functions	◐	In both display concepts we show more or detail contents to the side that is modulated.
DG6	Reflect consistency in assigning functions to gestures across applications	◐	Apart from the application switching technique, mapping of functions to gestures is consistent across the applications for both display concepts.
DG7	Reflect orthogonality in associating one unique function to one gesture	◐	In each interaction concept, we associate resizing of the display to one unique action.

Table 4.1: Overview of design goals that influenced our novel device concepts presented in this chapter and the extent to which they are covered in the previous work. ○and ◐show whether the requirements have partially been addressed in the previous work or not, respectively.

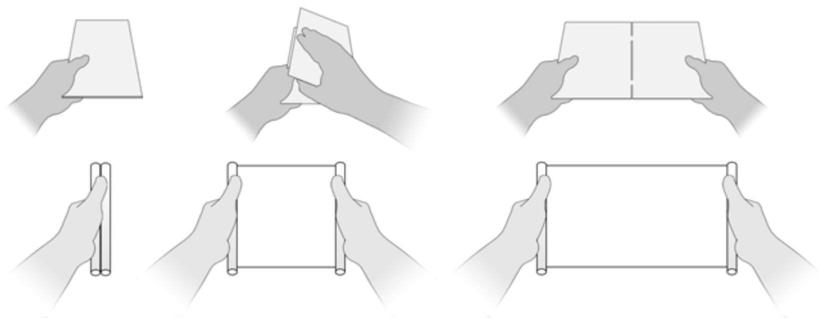

Figure 4.4: Resizing concepts: folding and rolling [Steimle 2012b]

Therefore, we want to systematically explore device concepts that their screen real state can be dynamically resized through *folding* and *rolling* techniques. Both techniques modulate the shape and size of displays similar to that of paper: they can get folded and become compact and small and when needed, one can expand them based on the current context of use. Folding gesture allows for deforming the display in discrete levels and stepwise manner whereas, rolling enables continuous and more malleable resizing of the screen real estate. While in the farther future, displays might ultimately become so flexible that they can be arbitrarily folded and rolled in any direction and along any axis, thin-film display technology is still quite far from this point. This is why, as a first step, we opted for unidirectional horizontal resizing for our device concepts. Another advantage of this is that upon a unidirectional fold or roll action, the rest form factor resembles that of current handheld computing devices and thus, we believe it may be more natural to users.

In this section, we present two novel device concepts: FoldMe – i.e., a foldable dual-sided display – and Xpaaand – i.e., a rollable display – along with their interaction design spaces. While touch is the most common input for today's hand-held computing devices and it is very likely that future devices featuring foldable or rollable displays will be equipped with (multi)touch input technology, we also look for plausible combinations of touch and deformation-based input for such devices. In our device concepts, we take a special care to design techniques that serve as an addition to touch interaction, not replace or impede it.

4.2.1 FoldMe: A Foldable Display Concept

FoldMe is a device concept that features thin-film, double-sided (front and reverse sides) displays that can be folded using predefined hinges. The device concept enables users to dynamically alter both size and shape of the display and also to access the backside using folding technique. Since folding is a complex technique with many variations, as a first step and for the sake of simplicity, we relaxed the device configuration so that its display can be folded along predefined hinges. The basic interaction units of such devices presented next are primarily organized based on different types and forms of display folds, followed by illustrating various touch-and-fold combinations.

4.2.1.1 Types of Folding

Fold-to-front and Fold-to-back

We consider the general case in which the hinge allows for fully rotating the flaps – i.e., they can define any angle from 0 to 360 degrees. In this case, each of the flaps can be folded toward or away from the user. In origami art, based on the form of crease, these are named as valley and mountain folds. These terms are well-suited for describing the actual state of the fold, but they do not account for the different directions involved in the folding process (fold or unfold). This is why we introduce the following terminology: fold- to-front and front-unfold for the valley fold; and fold-to-back and back-unfold for the mountain fold. These are depicted in figure 4.5.

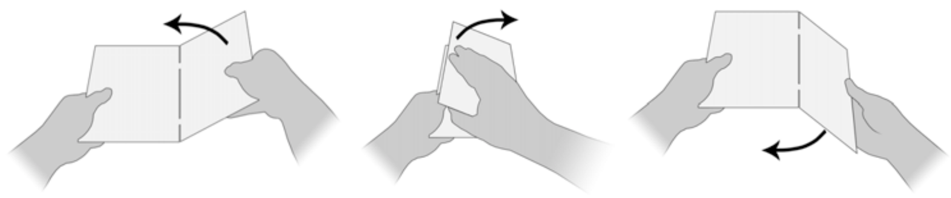

(a) Fold-to-front (b) Front-unfold (c) Fold-to-back

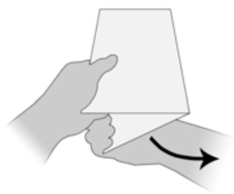

(d) Back-unfold

Figure 4.5: Basic folding gestures

Both fold gestures change the size of the available screen real estate. Yet, fold-to-front brings a portion of the reverse side to the front. In fold-to-back, it is vice versa – i.e., a portion of the front display is brought to the reverse side. These two types of folding are similar to the way we might handle a book or magazine. Folding over and front-unfold are used to open and close the book. Fold away allows for more convenient holding of the book while reading.

Continuous vs. discrete fold

Folding can be performed as a continuous rather than discrete action. In this way, starting from neutral state (flat state), turning the flap either toward to or away from the user results in a continuous input with positive or negative values (cf. figure 4.6a). This can be used for controlling continuous parameters, for instance.

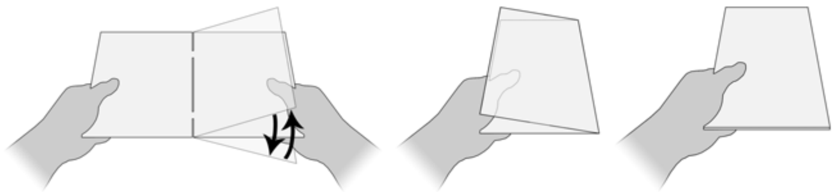

(a) Continuous folding (b) Temporal fold (c) Complete fold

Figure 4.6: Special types of folding

Temporal vs. complete fold

With physical books it is a common practice to temporarily index a page by putting the forefinger or thumb at that page. This enables a quick and convenient way for referring to a specific page while flipping through the pages of the book. Inspired by this practice, in our device design, we define the temporal fold as a type of folding in which one finger is put in-between the flaps as illustrated in figure 4.6b. In contrast, a complete fold requires that the flaps be placed directly on top of each other (see figure 4.6c).

4.2.1.2 Forms of Folding

The folding technique, yet simple, offers a rich set of alternatives. In designing FoldMe, we consider device configurations with displays oriented along a single axis (longer edge) in three different forms:

Centerfold Folding the display along its center axis to create two equally large flaps (cf. figure 4.7 a). This resembles a book metaphor. This form factor is similar to some existing devices, however, having display only on the front side. FoldMe offers displays on the reverse side too which is accessible by a folding technique.

Partial fold Folding the display along an asymmetric axis. In this case, folding-over the flap does not cover the entire screen (cf. figure 4.7 b).

Compound folds Different configurations can also be achieved by tessellating more display segments (n) with hinges (n - 1) in one dimension. We limit our design to two hinges with two equal sized display segments (cf. Figure 4.7).

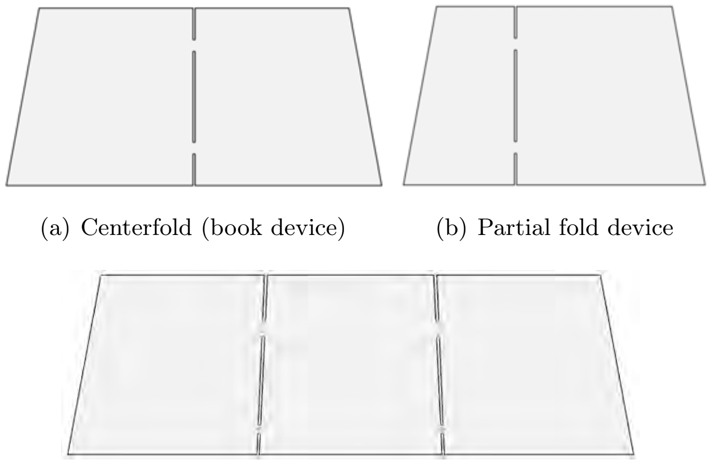

(a) Centerfold (book device) (b) Partial fold device

(c) Triple device

Figure 4.7: Various device configurations of FoldMe

We believe that study of these basic forms of folding, yet simple and restricted to one dimension, provides the foundation that is necessary for the design of more faceted future foldable devices. Moreover, these basic fold gestures maintain common rectangular form factors of today's hand-held devices that we are already accustomed to use. We leave more complex, diagonal, and origami-style folding interactions for future work.

4.2.1.3 Folding and Touch Input

We base our design on whether folding and touching are both performed with the same hand or whether they are delegated to different hands yet both occur at the same interaction cycle. In the one-handed case, the user holds the device with one hand while the other hand touches while folding. In the bimanual case, one hand touches while the other hand folds. These two cases are depicted in figure 4.8.

There are some cases in which it is difficult to touch while folding (e.g., touching the front page of the book device while doing a front-unfold). In these cases both activities are performed in a sequential order: first touch, then fold. No matter how

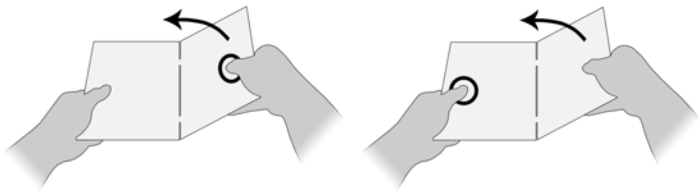

(a) (b)

Figure 4.8: Touch and fold combination performed either single-handedly or bimanually

it performs (in one interaction cycle or sequential), direct touch input on the display has the potential to distinguish between different commands that are all associated to one folding gesture and thus enabling a wider set of commands to be triggered by users.

One apparent issue in integration of touch interface is that users might trigger false touch events by simply holding the device. We propose two solutions for this issue. First, a bezel can assist users in preventing triggering unwanted touch events, similar to the bezel of current rigid handheld devices. In our device concept, however, this would require that each of the flaps is entirely surrounded by bezel, resulting in a very large bezel in the middle of the display when it is in an unfolded state. Hence, this approach on its own might not be a practical solution. Another solution is by automatically distinguishing between contact points that are created by the grip and actual touch events [Song 2011]. We believe that a combination of both approaches will allow users to hold the device comfortably while folding and touching.

4.2.2 Xpaaand: A Rollable Display Concept

We envision an ideal Xpaaand device that features a thin, lightweight, and high-resolution display that the user can (un)roll for dynamically resizing the display. Two handles on two opposing sides of the display provide for holding the device in a comfortable manner either in landscape or in portrait with one or two hands. The device hence the display cab physically expanded in one dimension by pulling both handles apart. It can be collapsed by pushing them together. Once the device is resized it maintains its new form factor.

By embedding, for instance, accelerometers in both handles, the device can sense which side is actuated. We distinguish three types of resizing: Pulling out or pushing in the left side of the display, the right side, or both sides (symmetric) as shown in figure 4.9. These three types can be mapped to different functionality.

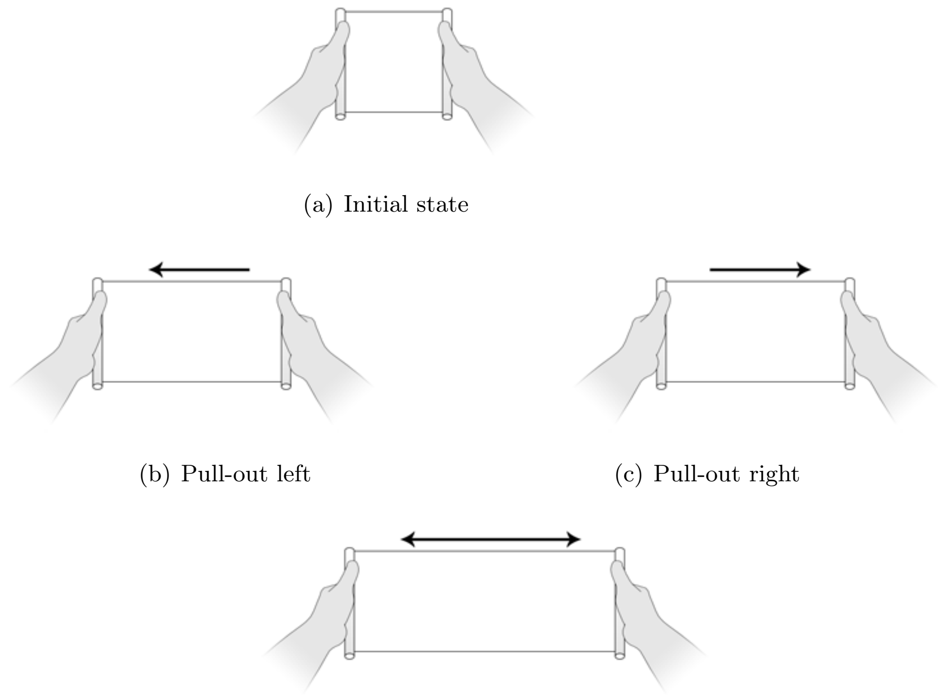

(a) Initial state

(b) Pull-out left

(c) Pull-out right

(d) Pull-out symmetric

Figure 4.9: Various types of resizing offered in Xpaaand

4.2.2.1 Metaphors for Rollable Displays

Figure 4.10 shows a traditional scroll that its content can be rolled in or out from one side (in this case left side). Depending on which side is pulled out, two distinct behaviors can be observed: by pulling out to the right side, content on the paper moves also with and it looks like that is being *locked-in-hand*. Pulling out the opposite side, however, resembles that content is remained attached to a fix position in the view point of users. We call this content *locked-in-viewport* metaphor. These two metaphors are also similar in nature to those proposed by Song et al. [Song 2009].

We consider the content locked-in-viewport metaphor to be better suited in most cases, as it is more natural and allows users to increase the display side to that they want to see more content. However, in some specific cases which we will discuss later in the next section, content locked in hand is more adequate. Therefore, in our design we support both metaphors consistently for interactions on both sides. In Xpaaand, when users simply pull either side in or out, we use content locked-in-viewport as the default metaphor and show more content to the side that is being manipulated. The users can, however, switch to locked-in-hand metaphor by using for instance, touch input (discussed below) or embedded physical buttons on each handle. Pressing the button on either side, means that the user can lock the content, which is displayed at that side, in his hand.

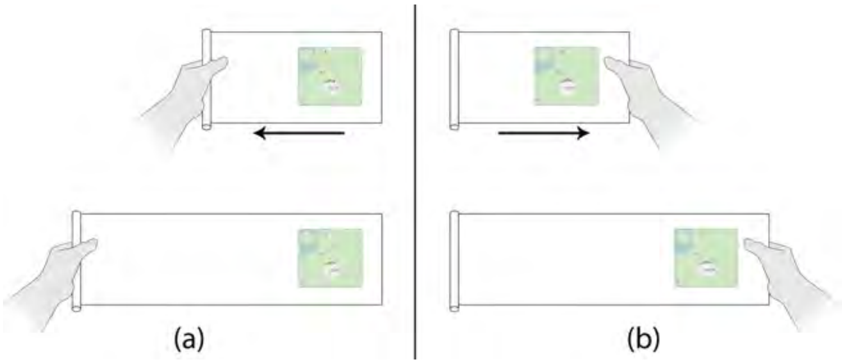

Figure 4.10: Metaphors of rollable displays: (a) locked-in-viewport and (b) locked-in-hand

4.2.2.2 Rolling and Touch Input

Similar to FoldMe, we consider the combination of rolling in or out with touch input that occur at one interaction cycle and whether it is performed with one or two hands. Thus, three main variations can be considered as illustrated in figure 4.11: resizing either left or right handles while touching or performing a symmetrical resizing (i.e., pulling both handles at the same time) while touching down. The latter resembles a sort of physical stretching of the content.

4.3 Interface Concepts

In this section, we propose novel interface concepts and interaction techniques for the resizable device concepts introduced above. We advocate dynamic modulation of screen real estate for accomplishing basic and common tasks found in the current mobile user interfaces. In the following, we first discuss the general interaction principles that provide basis for a set of more advanced interaction techniques presented later in this section.

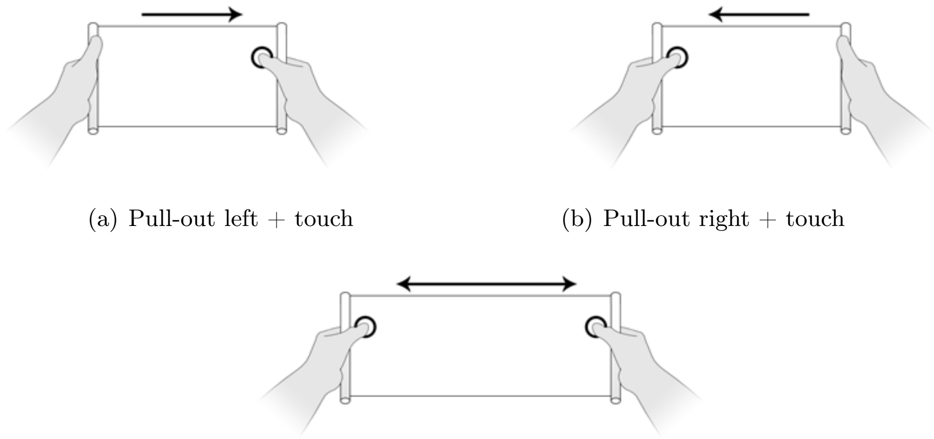

(a) Pull-out left + touch

(b) Pull-out right + touch

(c) Pull-out symmetric + touch

Figure 4.11: Combinations of resizing and touch input in Xpaaand

4.3.1 Interaction Principles

In general, both device concepts allow users to dynamically expand or shrink the size of an existing view, yet through different physical gestures. Based on the prior studies [Lee 2010] as well as inspired by real-world examples of paper-based artifacts – such as scrolls, maps, newspapers, books, etc. –, we propose a set of interaction principle for expanding and shrinking the display size as follows:

Expanding the display size

Regardless of how devices get expanded, it increases the available screen real estate that can be used for:

- *Displaying more contents:* The larger display size can be utilized to collate or extend current view, facilitating comparison of several items, and improves the display of contents that are better suited for landscape mode.

- *Displaying detailed contents:* The obtained screen space can be leveraged to display focused or contextual views of a dataset [Cockburn 2009]. For example, overview+detail interface designs can particularly benefit from the expanded view to show an overview of a spatially separated display segment generated on-the-fly. Moreover, users can physically activate the overview whenever is needed by expanding display or unfolding the corresponding display segment. Another category of interfaces that can take advantage of the expanded view is zoomable interfaces [Bederson 1996], which involve a temporal separation between views. We believe that overview+detail interfaces are more suitable for types of discrete expanding the screen real estate offered in FoldMe. On the contrary, the continuous resizing of rollable displays corresponds to the tempo-spatial nature of zoomable interfaces.

- *Exposing application functionalities:* Dynamic increasing of screen size through unfolding or unrolling can be used to trigger functionalities within or between application. As a concrete example, expanding can result in exposing different functions of an application or showing a different application.

Shrinking the display size

Apart from reverse actions of the aforementioned principle upon shrinking the display, we propose a set of other principles based on collapsing the display size. Although they are preliminarily designed for devices featuring dual-sided foldable displays, they can be used in designing interaction techniques for rollable displays. In previous work, shrinking the display through folding has been mapped to the *close* and *open* commands [Lee 2010]. In our design, folding provides access to the backside of display. The combination of folding action and the screen space on the backside opens up a set of novel actions:

- *Level up:* the user can have access to one level higher in a hierarchy using folding or rolling the display in. For example, in a photo browsing application, the user can fold over or roll in and expose the view of albums.
- *Overlay:* naturally, fold covers the front side of the display with a display layer coming from the backside. This can be interpreted as an overlay or an augmentation to the current view.
- *Accessing backside content:* in our design we also assign fold-to-front to access content or functionality on the backside of display. This is particularly designed for the FoldMe device concepts

4.3.2 Interaction Concepts

Based on the interaction principles we designed and implemented a set of more advanced physical interaction concepts to support basic tasks. Moreover, we present some interactions that make use of direct touch input while rolling or folding.

Physical Viewport Resizing

Regardless of type resizing (folding or rolling), expanding or collapsing the display primarily modulate the size of the viewport. Depending on which side of the device is manipulated, resizable displays allows for displaying additional content appeared or disappeared either on the left, the right, or both sides of the current viewport. While resizing rollable displays might physically alleviate navigating in a continuous digital document – such as a text, pictures, or a map that are scrollable, discreet folding action is better suited for displaying additional GUI components, such as contextual menus or application widgets. We have implemented three simple viewer applications: photo, map, and text (in portrait mode) running on Xpaaand. Figure 4.12 shows the photo application in which users can access additional contents by pulling out or in either handles. As another example, in the map viewer users can instead of panning, resize the display to see more content on the left or right side of the viewport.

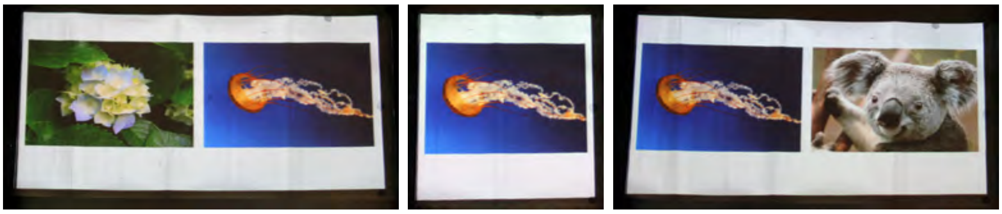

(a) Pull-out left (b) Initial state (c) Pull-out right

Figure 4.12: Photo viewer application showing the physical viewport resizing technique

To examine this technique on foldable displays, we developed a photo viewer application that separates controls and widgets from the main screen content (cf. figure 4.13). The flap of the partial-fold device acts as a physical tool palette while the bigger screen (primary screen) displays first-class contents. When needed, using back-unfold users can access the controls and widgets and once the interaction is completed, they are able to fold away the flap to the back of the primary screen. In this way users can remain focused to the main contents while reducing clutter on screen by off-loading interactive elements away from the display [Burstyn 2013]. Users also can interact with photos using common touch gestures: swiping to go to the next or previous images in an album. The tool palette displays several common controls and functions related to photo application, such as different buttons for sharing or widget controls for adjusting visual properties of a photo.

This interaction technique facilitates common activities, such as deeply focusing on a single item, comparing adjacent items, or getting a quick overview of many items. Furthermore, it enables a quick and intuitive accessing of controls and wid-

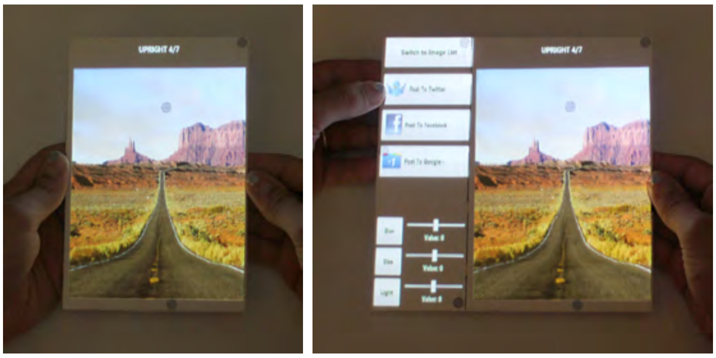

(a) (b)

Figure 4.13: Foldable tool palette

gets while keeping the primary screen free from any widgets and other distracting elements.

Physical Navigation in Menus

This technique supports the physical navigation through hierarchical lists (e.g., menus and folders). We utilize extending or shrinking the display to show or hide digital contents and the same time, trigger the open and close commands. In Xpaaand, to open a selected folder users pull the display out in the direction of the dominant hand. This selected item (e.g., the next menu level or contents of the folder) is displayed on the newly created display area, resulting in an overview+detail view. FoldMe also supports this technique by (un)folding the flap (the smaller display segment of the partial device). In order to display an overview of items of a folder users cab unfold the flap. To navigate even one level higher in the hierarchy, users can fold-to-front the flap on the primary screen. This displays an overview of different folders on the flap and a thumbnail view of items of the selected folder on the visible half of the primary screen. We implemented this technique in photo viewing application for both devices illustrated in figure 4.14.

Compared to today's fixed-size mobile devices, one important advantage of this technique is that through resizing the display, users observe both overview and detailed views of a hierarchy at the same time. In case of multiple levels, the technique retains the trajectory users traverse in the hierarchy. Due to the confined display size in rollable displays and limited fold numbers of foldable displays, this technique can support navigating in hierarchies with a few levels. Nevertheless, it can efficiently support going back and forth between levels that are frequently used by users.

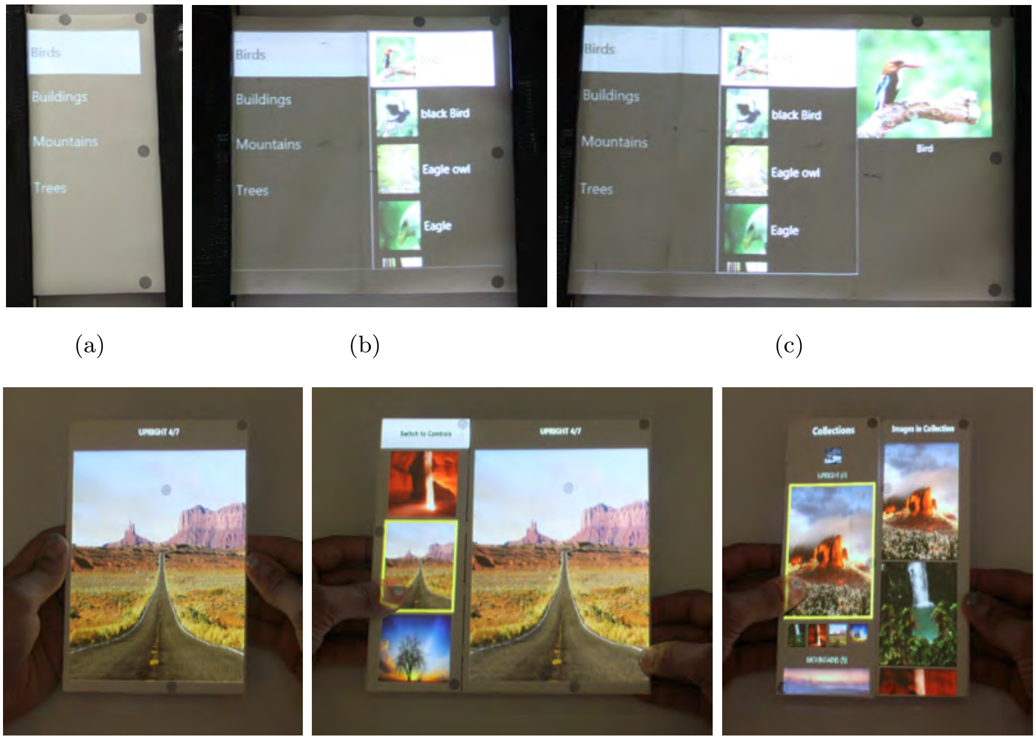

(a) (b) (c)

(d) (e) (f)

Figure 4.14: Physical navigation in menus

Physical Spinner Control

In this technique, we leverage the malleable nature of physical resizing for controlling continuous parameters. This means that instead of manipulating a spin control or a slider widget using direct-touch input, this technique leverages the continuous nature of folding or rolling for fine adjustment of values. It can be imagined as a tangible spin or slider control that increments or decrements a value when the display is resized. In this technique the main focus of the user, the primary screen, remains stable and well visible while the value is adjusted.

To translate the resizing input information (folding or rolling of displays) into parameter adjustment, two approaches can be identified: absolute and relative [Burstyn 2013]. In absolute mapping, resizing input is directly mapped to the state of a parameter. For example, in FoldMe, this means that folding input angle is directly translated into adjusting values: flat state of the display presents the zero value and folding towards and away from the users result in increasing or decreasing the value. The most and least values corresponds to +90 and -90 fold degrees. On the contrary, with the relative mapping, the resting state of display presents the current value of parameter, and any physical resizing results in value adjustment from the current value in the same direction. Due to the fact that our resizable display prototypes cannot retain the current display size (see prototype section 4.4) we used the absolute mapping approach for this technique.

(a) (b) (c)

Figure 4.15: Physical spinner control

This technique is implemented in the photo viewing application running on FoldMe. Users can adjust several parameters by partially folding the flap to the back or to the front. Spin controls allow for adjusting brightness, contrast and blur level of photos. Our implementation uses an absolute mapping between the folding angle and the amount of increment repeater. For each function, a button is displayed on the flap. The user can touch on one of the buttons and fold to adjust the corresponding value. Visual feedback about the current state is given by a simple slider on the right side of the spin control. This technique is illustrated in figure 4.15. To provide a natural polarity, we map folding the flap in positive (toward the user) or negative (away from the users) angles to increment or decrement the parameter.

The continuous pulling out the display can be exploited to support fine adjustment of a slider on devices featuring rollable displays. We sketched this interaction in video player application depicted in figure 4.16. Using this technique, pulling out the display can be seen as a physical instantiation of the rubber band metaphor used for temporal video browsing [Sun 2008]. Therefore, resizing of display is mapped to spanning a rubber band that is connected to the timeline of a video: the further the display is pulled-out, the stronger the band is spanned, the faster the video is played back. Pulling out to the right can be mapped to fast-forwarding, whereas pulling out to the left means rewinding, respectively. Physical resizing of display allows for occlusion-free and fine-grained interaction. On the other hand, the additional screen space can be used to scaffold a user's navigation process by allowing for a detailed overview over the key frames at the current playback position.

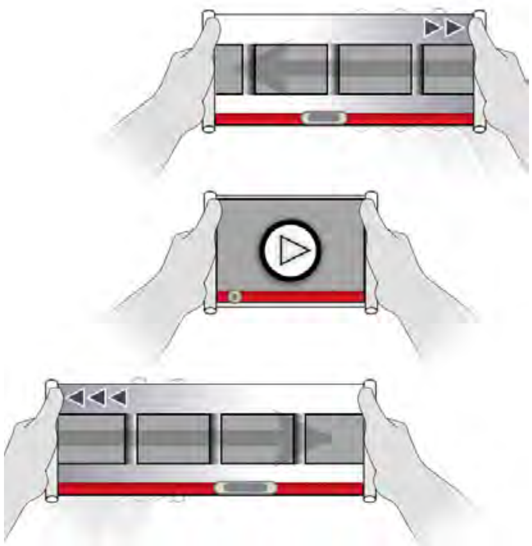

Figure 4.16: Physical timeline control

Physical Clipboard

This technique leverages dynamically created additional display space for temporarily holding an information element. This allows the user to compare items and creating a visual clipboard. By locking one information element in hand (pressing button) while expanding, the element is moved to the clipboard region, which appears on the fly. Moving the item back (e.g., for pasting) is performed with the reverse action. Collapsing the display without pressing a button hides the visual clipboard (locked-in-viewport metaphor). We have implemented an application for cut and paste with photos.

Physical Zooming

Zooming in or out is performed using the locked-in-hand metaphor. By pressing both buttons simultaneously, content gets locked-in-hand at both sides. While resizing the display, content on the viewport then gets stretched or shrunk, which is mapped to zooming. For navigating to very deep zoom levels, which cannot be reached by pulling out the display once, the user can iterate over zooming and viewport resizing several times (similar to lifting a mouse to move the device but not the cursor).

This technique reduces a problem of zooming on fixed-size screens, where magnifying the view results in losing context information. In contrast, enlarging the

(a) (b)

(c)

Figure 4.17: Physical zooming in maps

viewport while zooming in enables seeing the detailed view and still having more context visible in one dimension. We have implemented applications for zooming in 1D (i.e., semantic zooming within a text document, shown in figure 4.18) and 2D information spaces (i.e., geometric zoom within maps, shown figure 4.17).

We moreover designed a technique that provides for accessing additional information that is *embedded* within a document (e.g., a hyperlink within a web page). The user selects the hyperlink by touching it and simultaneously expands the display. The link target is expanded in place, within the context of the original document, leveraging the newly created display space. We have implemented an application for accessing additional images within a text document.

Physical Application Switching

One of the key features of today's handheld devices is multitasking: allowing users to quickly switch between a applications. With the iPhone or iPad, for instance, a common way of switching between applications is to double click the home button, which opens up a tray containing icons of all running applications. By clicking on an icon, the user can switch to an application. This approach, although simple,

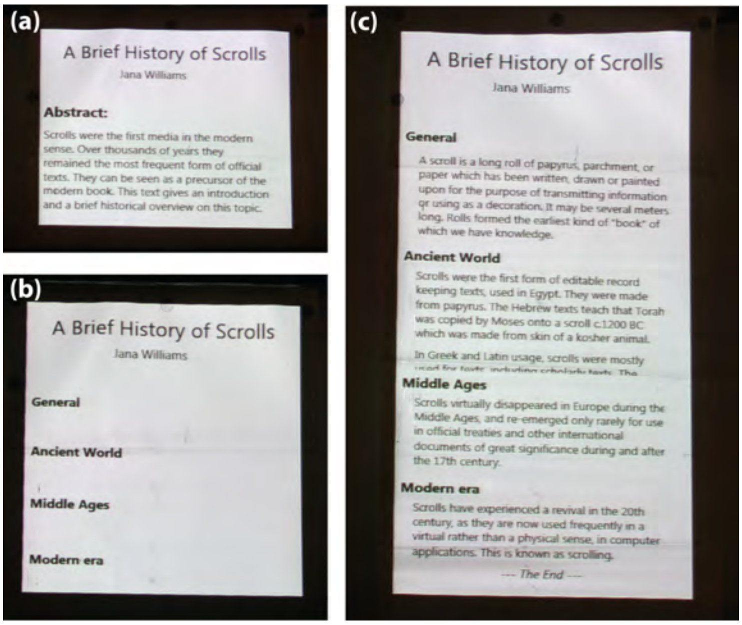

Figure 4.18: Physical semantic zooming

can become awkward in cases in which the user needs a quick means to frequently switch between several applications – for instance, when doing a text chat parallel to reading an e-book or when using Wikipedia for looking up definitions of terms that are used in an e-book. To mitigate this problem, we designed a technique for quick switching between applications that takes advantage of dynamic modulation of foldable and rollable displays.

One key affordance offered in FoldMe is that one can fold the display for accessing content that is located on the backside. This feature is particularly useful for an effective, embodied multitasking. Having one or several applications assigned to the backside, the user can quickly refer to them and then get back again to her context. Since each application is assigned to one unique page of the foldable device, rich spatial cues guide the user and, in combination with folding, generate a more direct and physical experience of multitasking. In this technique a complete fold results in closing the foreground application. In contrast, finger bookmarking keeps the application running in the background while the user is working with another application.

We also designed and implemented a physical application switching technique for rollable displays. This technique allows users to easily switch between a small number of frequently used applications by resizing the display to the size that is associated with a specific application. This is shown in figure 4.20. A similar ap-

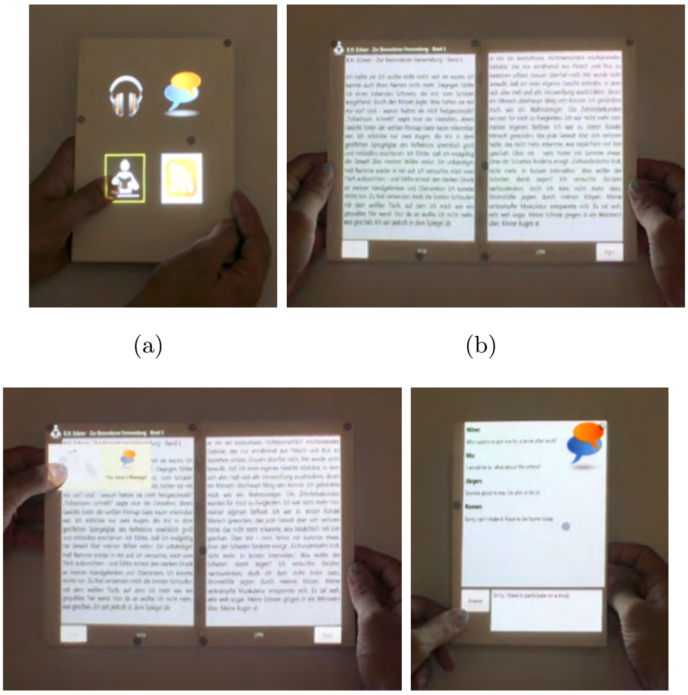

(a) (b)

(c) (d)

Figure 4.19: Physical application switching

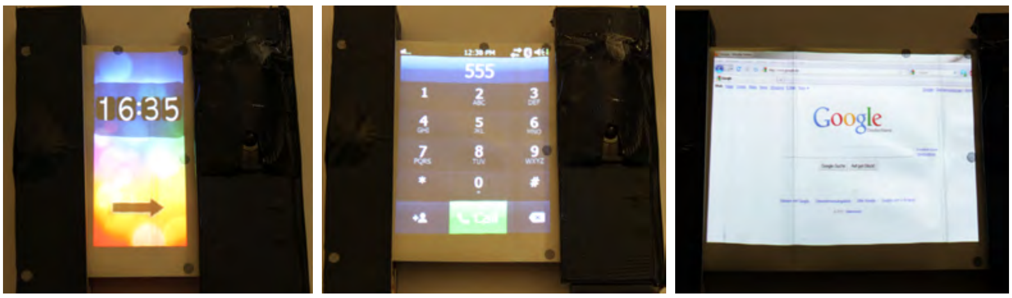

(a) (b) (c)

Figure 4.20: Physical application switching

proach can be taken for exposing or hiding functionality within one application.

Physical Layering of Information

This technique supports physically overlaying the information displayed on the front screen with the back-side display. This is slightly different to the other techniques of FoldMe in the sense that folding does not result in a different view but rather augments the contents of the front display. This is similar to applying a lens or placing a display with a transparent background on it. Having a device with more than one hinge allows the user to combine or merge different layers on top of the front display.

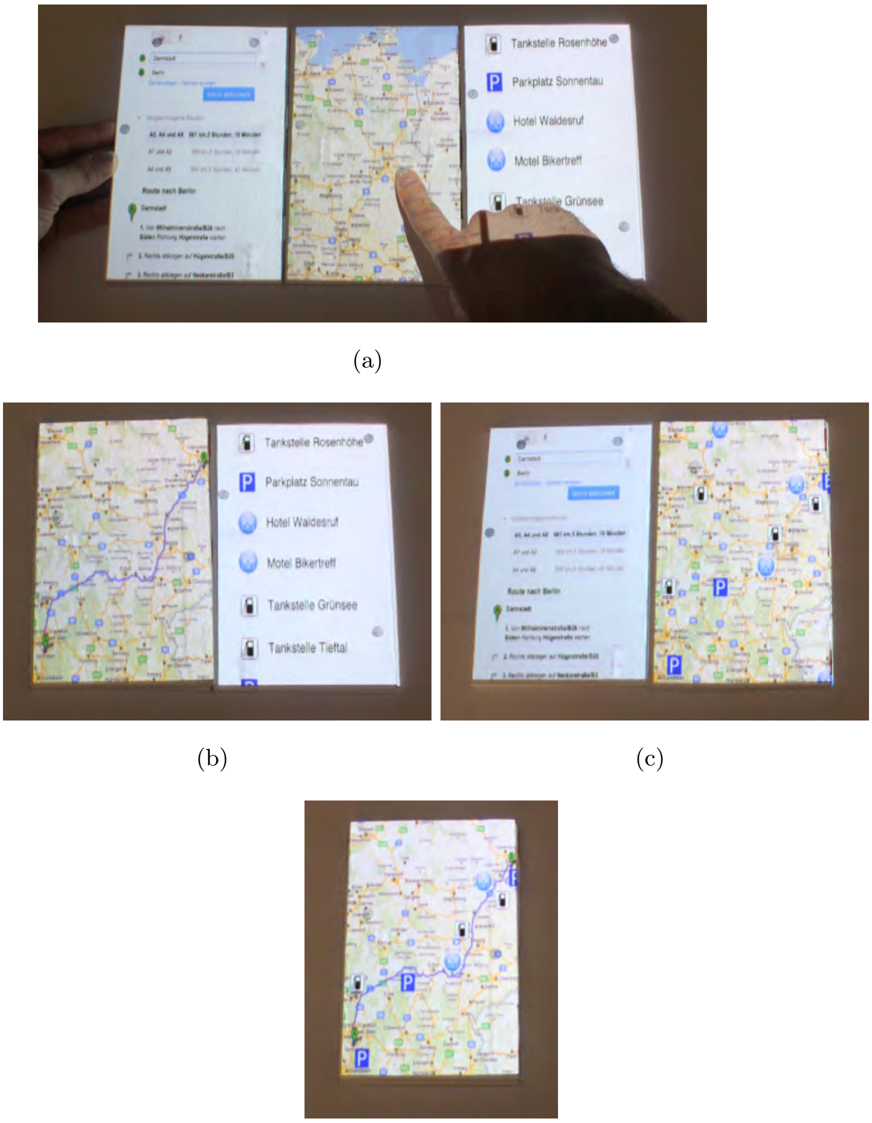

(a)

(b)

(c)

(d)

Figure 4.21: Physical layering of information through folding over

We implemented the technique in a map scenario running on a foldable device featuring three display segments (cf. figure 4.21). The application displays the Google map view as the primary view on the middle screen. Two adjacent screens represent textual explanation of a route (left screen) and a set of point of interests (POI) – such as gas stations, hotels and park places (right screen). To graphically show the route on the map, the user can fold the left screen with the root information over the map view. Similarly, the user can fold the right screen over the map view to augment the map with icons showing the location of POIs. In order to merge the two views, the user can fold both screens over the map. This results in a map view with integrated route and POIs (cf. figure 4.21d). Table 4.2 summarizes all the interaction concepts for both FoldMe and Xpaaand.

Underlying Principle	Description
Physical Viewport Resizing	
- Displaying more contents	Collating/Extending the viewport by pulling out or unfolding the display
Physical Navigation in Menus	
- Displaying more contents - Exposing application functionalities	Selecting a menu item while resizing the display results in navigating to one level deeper in a hierarchy.
Physical Spinner Control	
- Exposing application functionalities	The continuous action of resizing (through either folding or rolling the display) is mapped to adjusting an application parameter.
Physical Clipboard	
- Displaying more contents - Exposing application functionalities	Dragging out an item from a list while expanding the display triggers a cut-and-paste action and displays the item on the extended display area.
Physical Zooming	
- Displaying detailed contents - Exposing application functionalities	Expanding the display while touch-down on both sides of the display results in a zoom-in action.
Physical Application Switching	
- Exposing application functionalities	Resizing of display is mapped to switch between applications based on the display size and form factors.
Physical Layering of Information	
- Displaying more contents - Accessing backside contents - Overlay	Folding displays augments the primary preview with additional information. Folding more than one display merges information displayed on the primary preview.

Table 4.2: Summary of physical interaction concepts for resizable displays and their interrelation to the underlying interaction principles.

4.4 Implementation

Despite recent advances in creating thin-film display technology, currently available flexible displays do not yet allow for producing a device with rollable displays or very thin double-sided foldable displays that is untethered. Thinness of displays is crucial for free and easily folding and rolling in all directions. Moreover, current tethered flexible displays [Lahey 2011] limit the total thinness of a fully folded or rolled device. Therefore, we follow a passive display approach used in the previous systems [Holman 2005, Lee 2008, Steimle 2013], in which arbitrary flexible surfaces are augmented with a corrected top projection. This enables us to realize our device concepts and deploy our interaction concepts regardless of technological barriers. Since this approach closely emulates look and feel of actual foldable and rollable displays we are also able to test our interaction concepts through user studies.

In the following, we detail on physical prototypes of FoldMe and Xpaaand followed by introducing the simulation environment that is used for augmenting them.

4.4.1 Device Prototypes

4.4.1.1 Xpaaand Prototype

The Xpaaand prototype (cf. figure 4.22) consists of a physical scroll made of white foil, where display content is projected. It can be resized to widths from 5 to 39 cm at a fixed height of 18 cm. It is 4 cm thick and weighs about 900g. A box is attached at each of the two ends. These boxes act as physical handles for easily grasping the device. They contain the scroll and electronic components. A physical button is positioned at the center of each of the boxes. As our current prototype does not support touch input and to support navigation on the display while it is held with two hands, a trackball is integrated into one box. It can be manipulated with the dominant hand (left-handers can rotate the device by 180 degrees). Wireless communication links the prototype to a nearby PC that hosts tracking software and applications. To keep the device stable, straight, and stiff at different display sizes, we fixed an expandable strut on the back of device.

4.4.1.2 FoldMe Prototypes

The FoldMe prototypes incorporate the main physical properties of our display concept described in 4.2: lightweight, rigid, and easily foldable in both directions. Furthermore, they become stiff when in an unfolded flat state. This provides for a single unified display space that can be easily held from one side.

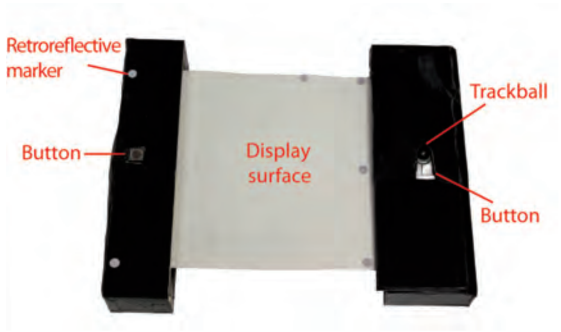

Figure 4.22: Xpaaand prototype

We used foamboard that is made up of a sheet of foam sandwiched between two sheets of cardstock paper. To connect two sheets of the foamboard, we embed a number of relatively strong magnets into the long edge of each sheet, which realize a snap-in effect when the display is in an unfolded flat state. In order to ensure that sheets cannot be detached and hinges remain stable and well-aligned during manipulation, we integrated three straps connected to both sheets. A schematic view of the prototype design is depicted in figure 4.23a. In this way, no gap is visible at the hinges. Users perceive only a slightly visible crease at the possible fold location. As it can be seen from figure 4.23b, we constructed three different prototypes – namely, book, partial-fold, and dual-fold devices. They allowed us to examine different types and sizes as well as more-than-one-fold variations.

4.4.2 Simulation Setup

Our simulation environment consists of an Optitrack motion capture system[1] with six infrared cameras, an array of full HD projectors mounted on the ceiling and several foldable prototypes augmented with infrared retro-reflective markers. The information provided by the tracker system (position, orientation, the state of the prototypes) is used to warp the projected image onto the prototype in real-time. Display contents are then projected by two full HD projectors. The display update rate is 60 fps; the average resolution on the display is approximately 42 dpi. A schematic overview of our simulation setup is depicted in figure 4.24.

In the simulation framework, we simulate the environment by constructing a Direct3D world model of Microsofts's DirectX[2] framework. In an initial calibration step based on the direct homography approach [Zhang 2010], we instantiate a Di-

[1] http://www.naturalpoint.com/optitrack/
[2] http://windows.microsoft.com/de-de/windows7/products/features/directx-11

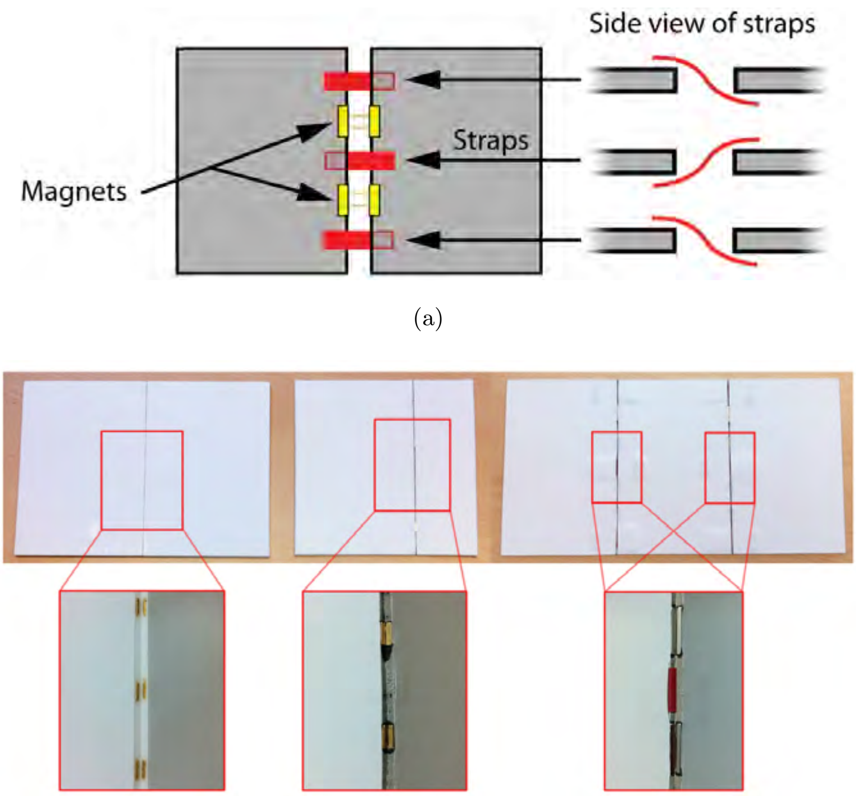

(a)

(b)

Figure 4.23: FoldMe physical prototypes: (a) schematic view of hinges using magnets and straps and (b) three prototypes from left book, partial-fold and dual-fold devices.

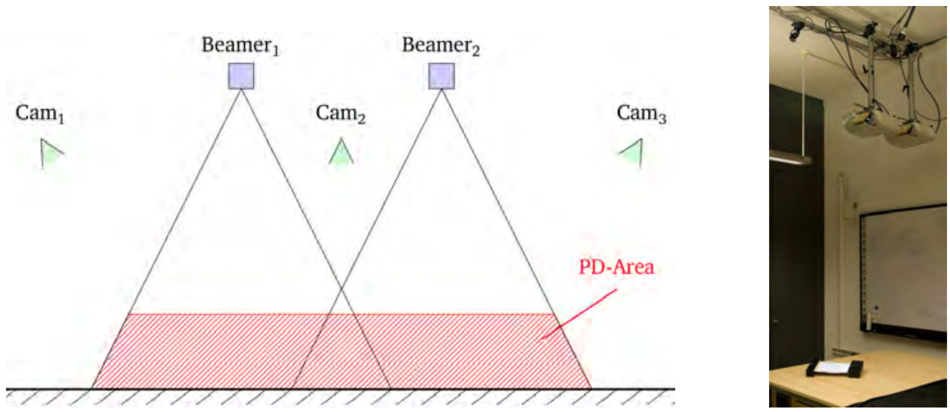

(a) (b)

Figure 4.24: Schematic view and hardware setup of the paper-like simulation environment

rect3D camera in the model with its position and orientation precisely set to that of the corresponding projector (an evaluation [Riemann 2010] showed that our calibration approach yields a precision of about less than 4 mm). Thus, the camera *sees* the passive display and its contents from the correct perspective. The camera view that is generated by Direct3D is displayed by the projector while the world model is continuously updated by the tracker data. In order to recognize different folding or rolling gestures, we implemented a gesture recognizer module that analyzes positional information of the display. The Direct3D model receives real-time screen captures of Windows Presentation Foundation[3] applications and renders them as textures onto the display.

In order to enable touch input, we attach an infrared reflective marker (IR marker) to one finger of each of the user's hands. Once the finger marker is sufficiently near to the display surface, we calculate the projection of the finger point onto the display plane using planar geometry. As this approach would fail when a touch occurs very close to one of the markers that are used for identifying the display, we took special care of positioning markers on the display surface in a way to ensure that the main screen area remains marker-free. Different approaches (e.g., by using pressure sensitive or capacitive touch foils[4]) were not possible since these either require tethering or too bulky electronic components.

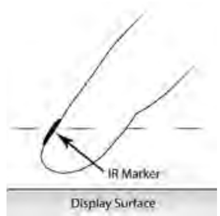

Figure 4.25: Touch implementation

The physical prototypes and the simulation environment allow us to closely emulate future devices featuring foldable and rollable displays. We implemented the interaction concepts introduced in 4.3 running on FoldMe and Xpaaand prototypes. This enable us to test our interaction concepts in user studies presented in the next section.

4.5 Evaluation

In this section, we evaluate the interaction concepts designed for resizable displays. This is done through two user studies in which we individually examined foldable and rollable device concepts. Our main goal of the studies was to obtain qualitative insights on how people generally interact with and apprehend physical resizing of displays as novel input modality for such devices. We aimed not to statistically

[3] http://msdn.microsoft.com/de-de/library/ms754130(v=vs.110).aspx
[4] http://www.sensibleui.com/

proof of a priori conjectures due to the high level of novelty offered in our design. Therefore, the studies illustrated a form of exploratory observations aiming to discover common interaction patterns and unexpected phenomena. In the following, we first report the FoldMe and then Xpaaand user studies.

4.5.1 Study1: FoldMe

In order to evaluate the FoldMe device and its interface concepts, we conducted an exploratory study. We were interested in observing how participants used our interaction concepts for identifying strengths and weaknesses, and also to see how easily users understood the novel interaction style of double-sided foldable displays. Moreover, we examined our device configurations and various functional roles that can be assigned to each folding type and form.

4.5.1.1 Participants and Tasks

We invited ten volunteer participants (nine male, one female) to our lab. Three of them were left-handed. All were professional computer scientists. Each participant owned a smart phone that they used on a daily basis. Five participants owned a general-purposed tablet for longer activities – for instance, reading and surfing. E-book reading devices were used by two participants.

The study consisted of two blocks. The first block focused on examining our device concepts and various configurations while there was not projection provided. We presented the paper prototypes and invited the users to think of an associated functionality for each device configuration and different scenarios in which they could imagine using the prototypes. No information was displayed on the device. However, we utilized some printouts of screen captures of some common mobile applications to foster brainstorming.

In the second block, we evaluate our interface concepts using high fidelity prototypes – i.e., augmenting our paper prototypes with interactive projections as described above. After a brief introduction, participants could explore each interface concept until they feel comfortable using the device and performing respective interaction techniques. Then, we asked them to search for a specific photos using our physical menu navigation technique and adjust its brightness using our physical spinner control. For the book device, we emulate a reading scenario using an eBook application running on both front displays of the device. On the reverse side of device, we assign a chat and music player application to left and right display segments. Users were asked to read the ebook and respond to a number of emulated notifications of the background applications using our physical application switching technique. We also test the physical overlay technique using the triple device and asked participants to search and show a specific POI along several local routs. Participants were asked to perform the tasks while thinking aloud. Sessions were video recorded for post analysis. The study was conducted in single user sessions that lasted approximately 90 minutes each.

4.5.1.2 Results

General Observations

All participants appreciated the concept of double-sided display and the possibility of folding and unfolding to modulate the screen size and shape. They were able to quickly understand the interaction principle of fold-to-front and fold-to-back in combination with the reverse side of the display. They found folding very natural (particularly the fold-to-front) and useful in terms of having a compact form factor for mobility and enlarging the screen when the context permits. One participant commented, "Folding [over] is kind of a shortcut to access the backside compared to flipping the whole page." All participants appreciated the strong sense of going back and forth that is suggested by folding. Folding was also found to be very appropriate for temporarily referring to something on the back side and then returning back to the front side. Almost all participants expressed privacy concerns of having contents displayed on the back side. They mentioned that it has to be inactive while working on the front side.

In the following, we explain results of the first and second study blocks in turn.

Device-Related Results

As mentioned before, in the first block of the study, no information was provided on devices. This allows us to gather qualitative data about each device and functional roles of displays. We categorized the results based on the devices presented as follows:

Book device: All participants commented that the book device is suitable for active reading scenarios. Fold-to-front and front-unfold were found to resemble close and open actions. We observed that since all four displays have the same physical appearance in size and shape, users did not map different roles to the displays but treated all of them as having the same role. P3 mentioned, "Having multiple displays on one side is useful for dividing your coupled task [like reading and writing] and the back side can be used for background applications that you need to refer to from time to time." Four participants explicitly mentioned that it might be suitable for occasionally referring to the back side.

Six participants found fold-to-front to be more intuitive and practical than fold-to-back, particularly when using the book device. One participant commented that "fold-to-front [from right to left] is like clicking the right mouse button to access the context menu, which is missing in touch interfaces" and is like *a special function.* Another commented, "Fold-to-front to access the back side to see more information or invoke other applications is very intuitive."

Partial-fold device: Users reported that the smaller flap of the partial-fold prototype is an *add-on* to the main screen and offers additional functionality, similar to a tool palette. P1 stated, "It [the flap of the partial-fold device] is very quickly accessible and easily foldable. It is more convenient than folding using the book prototype." Another participant commented, "Out of sight, out of mind but still quickly accessible." Six participants stated that the flap of the partial-fold prototype is suitable for having a list (or an overview) while the main screen displays the detail.

Dual-fold device: Seven participants perceived the dual-fold device as an extension of the book device. In general, participants had difficulties in imagining an appropriate application of the dual-fold device. However, they could envisage using the device for specific purposes – such as a graphics editing program or a map application. Almost all participants expressed privacy concerns of having contents displayed on the back side. They mentioned that it has to be inactive while working on the front side.

Interface-Related Results

In the second block of the study, we examine our interaction concepts running on devices. We report our findings based on each technique in the following.

Physical application switching: was very well received by all participants, particularly the combination with finger bookmarking that allows for *haptic feeling of pausing your context* and then unfolding to resume. They appreciated using the back side display for accessing the background application, and the front display for the foreground application. Three participants emphasized on the fact that it is suitable for less frequent and temporal application switching. In cases where tasks are highly interwoven (such as reading and simultaneous note taking), participants preferred to have both applications on the foreground display divided on two adjacent screens to avoid the need for highly repetitive folding.

Physical navigation in menus: Although the technique for changing levels in a hierarchy with the fold-to-front gesture was limited to two levels, it was well received by six participants in the partial-fold device. The technique provides a rich physical experience of going up or down in the hierarchy. However, in cases that the levels of hierarchy were more than two, participants preferred alternative ways of changing levels (e.g., direct touch input) and using folding only as a shortcut for switching between two frequently used levels.

Physical spinner control: This technique received positive feedback for several reasons. Participants commented that it supports interacting with digital content while both hands are holding the device and avoids hand occlusion on the primary screen while manipulating the content. However, three participants were unsure about the effectiveness of this technique in terms of efficiency and accuracy in comparison to direct touch interaction with a slider control. One reason for this might be the fact that our implementation required participants to interact with the spin control using their non-dominant hand. Another possible shortcoming might be the direct mapping of the degree of folding to increment of value.

Physical integration of information: Initially, this technique on the dual-fold prototype was found to be difficult to understand by most of the participants, since it is conceptually very different from existing mobile devices. However, after a short while they could figure out the idea behind it. Two participants saw a natural mapping to *do* and *undo* actions. P3 stated, "When I perform [fold and unfold] gestures to quickly check if a filter is appropriate or not, it resembles [do and] undo actions." All participants appreciated having the middle screen as the main preview of the application (in this case the map) while folding the left or right sides over preserves and only augments the main view, even though the physical screen size is changed.

The discussion with participants revealed that having the primary preview on one physical display and the filtered preview on the back side of another display gives them a safe feeling that the original content remains unchanged. Therefore participants frequently folded and unfolded the device repetitively, switching back and forth between different views. In contrast, results of the foldable multitasking technique showed that where the fold-to-front gesture entirely changes the current

General Findings
– Folding (particularly fold-to-front) was perceived to be very natural and quick.
– Compared to flipping the whole display, folding provides a shortcut for accessing the back side.
– Users raised privacy concerned about the content to be displayed on the back side of display.
Device-Related Findings
– In book device fold to front perceived as special function for invoking other applications.
– Book device affords for physically delineating applications to each display segment.
– Folding the flap in partial device was found to be easily and quickly accessible.
– There was a strong agreement that dual-fold device is difficult to use due to large folding possibilities it offers.
Interface-Related Findings
– Physical application switching was found to be suitable for less frequent and temporal task switching.
– Physical navigation was used for frequently used levels in a hierarchy.
– Users were skeptical about the practicability of physical spinner control.
– The Physical overlay technique resembled a do-and-undo action.

Table 4.3: Summary of the main findings from the FoldMe evaluation

view, users tend to perform the fold gesture more selectively and infrequently.

Table 4.3 summarizes the main findings of the FoldMe study.

4.5.2 Study2: Xpaaand

To evaluate our concepts of rollable displays, we conducted a second explorative study to examine how easy it is for novice users to understand the novel interaction style of Xpaaand.

4.5.2.1 Participants and Tasks

We recruited eleven volunteer participants, from which five were female, with an average age of 26 years. Three were left-handed. The study was conducted in our lab environment. We used the Xpaaand prototype described in 4.4 augmented with interactive top projection. The study was conducted in single user sessions that each lasted approximately one hour.

After five minute of introduction, participants explored the interaction concepts and were asked to perform the following tasks in the horizontal mode: searching and comparing photos using physical viewport resizing, finding certain photos in different albums using the physical navigation technique, performing multiple times the cut and paste operation with the aim of the physical clipboard, and performing physical application switching while thinking aloud. In the vertical mode, participants had to perform a reading task with and without the semantic zooming technique, followed by revealing a hidden figure using the accessing additional information technique. The session ended with a semi-structured interview. Sessions were video recorded for post analysis.

4.5.2.2 Results

All 11 participants appreciated the possibility of changing the display size dynamically. Physical expanding and collapsing of the display for interacting with contents was found to be *intuitive* and *enjoyable*. P3 commented, “It is great and makes so much fun because it imitates something physical and has haptic feedback.” P4 stated, “I have just done it, without thinking, it feels more active to me.” Although “it is really unusual to pull out” (P2), all participants quickly got accustomed to the resizing technique and performed it with ease and confidence after a few minutes. The locked-in-viewport metaphor used for viewport resizing was appreciated by all participants. One participant (P2) commented, “It is so intuitive because the contents remain entirely unchanged. The rest is just temporarily hidden in the scroll.”

Almost all users (10) reported that dynamic and continuous resizing helped them in searching and comparing items in a list. For a search task in the photo application, seven users intuitively enlarged the viewport to gain a broader view before scrolling through the list. Moreover, eight participants explicitly stated that the comparison of items was much easier in the flexible display size. Nine users enlarged the viewport before comparing adjacent items. This helped them to directly engage in more than one item while maintaining the full resolution of each item. P3 commented, “I intuitively resized it to a size which allowed viewing two photos one besides the other.” Even though users had enough physical space for expanding the display to its full size, we observed some instances of collapsing the display to focus on a specific location of the viewport. Furthermore, during the interviews three users explicitly mentioned to envision that they can effectively use the device in a restricted space very much like one handles physical maps or newspapers. These results indicate that intuitive resizing of the display eases search, comparison and serendipitous discovery of contents.

Besides physical viewport resizing, the physical interaction concepts for zooming and hierarchy navigation were particularly well received. The participants appreciated the combination of resizing the display and basic interface operations (e.g., open, close, zoom) which provide for a rich physical experience of going in and out (seven users) and keeping the context visible (three users). Participants reacted very positively to the technique for temporarily holding an element, which eased visual comparison. Physical application switching running on Xpaaand was found simple, but most participants needed one to two minutes to get accustomed to this technique. In contrast to other techniques, this was because in this technique, resizing displays resulted in invoking an application that does not use any of the metaphors of physical scrolling.

Most of the time users held the device easily with their two hands, their arms placed on the table. However, seven users felt that the prototype was too heavy for comfort, particularly for holding in a vertical orientation when widely expanded. Five participants had concerns related to physical fatigue. Physical resizing "is great but [it] might become a sport doing it too often or for a long time" (P5). With one single exception in the photo application, all these comments were raised during navigation in hierarchies when users desired to repetitively and quickly switch between several folders. The physical effort of repeatedly collapsing the device for going up one level and expanding again was considered too high, particularly because the device eventually had the same size as before. Several users stated that in this specific case, they prefer having alternatives (e.g., touch input) to the proposed technique.

Finally, as an unexpected advantage collaborative uses of the device were suggested. This comprised temporally enlarging the display for showing photos to other people (P2, P10) or games (P9) where each of two users holds the display at one side.

Table 4.4 summarizes the main findings of the Xpaaand study.

4.5.3 Discussion

The results of both studies showed that dynamic modulation of display is strongly preferred. Participants commented that physical resizing of display has potential

General Findings
– Physical expanding and collapsing the display were found to be intuitive and enjoyable.
– Dynamic resizing assisted searching and comparing items in a list.
– Rolling out the display was interpreted as offering options, whereas pulling in the display resembles committing changes.
– Physical zooming and menu navigation resembled a sense of going in and out into contents.
– Repeatedly collapsing and enlarging the display to accomplish a task was not preferred.

Table 4.4: Summary of the main findings from the Xpaaand evaluation

to increase the user experience and effectiveness of interaction with digital contents on hand-held devices. From participants' comments, we found out that physical resizing as an input modality is *natural* and *intuitive*. It is *obvious how it works*, and it requires *less concentration*.

Through the studies, we also identified one major concern associated with physical resizing of display as an input modality – i.e., the *physical fatigue*. Compared to direct-touch input, which is less tiring, physical resizing involves both the hands and the arms in the interaction. We observed that participants were concerned that physical resizing will become a *sport* while performing repetitive or frequent commands. This issue signifies the need for an appropriate trade-off between user experience, efficiency, and physical effort. We believe that this limitation will become less severe with future devices that are more lightweight than our prototype and provide alternative input techniques, such as direct-touch input.

We observed that the FoldMe device was treated as a *multitask system* (particularly in the book device and the dual-fold device) due to the multiple equal display segments it offers. There was a strong agreement among participants that distributing applications across different displays while being able to access them concurrently is the prominent feature of the FoldMe device. We conclude that the FoldMe design mimics the kind of tactile-kinesthetic affordances of physical paper and offered a rich multi-display experience. This is inline with the findings of [Chen 2008, Hinckley 2009], which revealed similar affordances of dual-sided devices. It was also found that the front display segments are suitable for frequently used applications and, thus, support highly interwoven activities, whereas back side of displays are for background applications. In contrast, findings of the Xpaaand study showed that due to the fact that rollable displays offer one unified surface, it was found to be more effective while users are engaged in one application. Thus, interaction concepts in which resizing the display resulted in collating the viewport, zooming of data objects, or showing additional contents in one application were mainly preferred.

One-dimensional folding and rolling (horizontal) in both device concepts as well as predefined hinges in the FoldMe design prevent us from investigating free-form folding and rolling variations. We believe that the techniques and results of the studies can be extended to fully flexible displays. Fully flexible displays will allow the user to select the position and the extent to which the display gets folded or rolled. This will render possibly more flexible sizes of folded application windows, tool palettes, and lenses.

4.6 Conclusion

Today's computer user interfaces do not leverage many of the possible degrees of freedom offered by everyday flexible materials such as paper. Given the rapid advances in thin-film display technology, such as E Ink and organic light emitting diodes (OLED), we will witness a radical change in the design of computing devices. These technologies potentially provide thin, lightweight and even deformable displays that incorporate many of the physical properties that until now were unique to paper documents. Multitouch displays might ultimately become so thin that they can be arbitrarily folded and rolled while featuring high-resolution display both on the front and the reverse sides. In this chapter, we investigate how future flexible displays might adopt dynamic resizing of display through folding and rolling techniques. We particularly exploit the dynamic modulation of screen real estate for interaction with digital contents.

Based on literature analysis on deformation-based and flexible-display user interfaces, considerable research has focused on bending and twisting gestures. While informative and practical, flexible displays' promise goes beyond such subtle deformations of curving the entire or part of the display. Through contributing two novel device concepts – namely, FoldMe and Xpaaand – we explore more a paper-like deformation of display that is based on folding and rolling techniques. The device concepts can potentially provide a solution to the traditional dilemma of fixed-size displays: increased screen size vs. portability. This means that instead of navigating through content spanning outside of the display area, users can physically extend the display screen real estate of foldable or rollable displays to collate their viewport to see more content.

We moreover leverage the physical modulation of the screen to provide an additional and intuitive input modality. Based on the natural way of interacting with paper-based artifacts – such as paper rolls or newspapers – we explore physical interaction design spaces of foldable and rollable displays. We characterize the design space of foldable displays based on various types, numbers and forms of folding. Inspired by traditional paper rolls, we also defined two different metaphors for rollable displays: locked-in-hand (i.e., when the content moves as users pull out the display) and locked-in-viewport (i.e., content is attached to the display and do not move as users resize the display). In addition, we looked at the possible combination of incorporating the direct-touch input while folding and rolling displays.

In the present chapter, we also contribute user interface concepts that make use of the full potential of the physical resizing input modality. We first define several interaction principles that serve as the foundation for designing concrete interaction concepts. The principles basically include different possibilities that resizing the display might lead to: seeing more content, or more detailed content, as well as invoking a command. Due to the larger variations in FoldMe, folding over the device can additionally be interpreted as going a level up, overlaying the current view, and accessing content on the back of the display. Based on the interaction principle, we design a number of physical interaction concepts to accomplish basic tasks – like navigating, zooming, controlling parameters, and application switching.

Following up a passive display approach, we implement the interaction concepts running on prototypical realization of foldable and rollable devices. To do so, we build a simulation environment to emulate paper-like displays in which passive flexible surfaces are augmented with a top projection. This makes it possible to evaluate our interaction designs by conducting real user studies regardless of barriers in using actual foldable or rollable display technology, which is still in its infancy. We conduct two exploratory user studies to test the concepts of rollable and foldable devices individually. Results of both studies show that physical resizing can considerably improve the effectiveness and the user experience of interacting with digital contents on hand-held computing devices. We observed that since FoldMe offers physically separated display segments, it was perceived rather as a multitasking system and provided for a rich multi-display experience. Xpaaand, however, enabled users to physically engage in data objects of an application on a unified yet dynamically resizable display. Results also suggest that in situations where users have to repetitively resize the display, alternative less tiring input modalities, such as direct-touch, are more preferred.

Chapter 5

Conclusions

Contents

Knowledge workers develop highly efficient practices of manipulating physical objects during a liftime of working in real desktop workspaces. A set of these practices involves organizing objects found on the surface of tables – such as moving, shuffling around, or laying out. In addition to these table-centric practices, knowledge workers also develop a high level of manual dexterity and skills with respect to the category-specific manipulation of physical objects; this is particularly true for paper, an indispensable component of information-based activities. Paper-specific techniques include piling or stacking of physical documents for organizing the desktop workspace [Malone 1983, Mander 1992], as well as flipping, folding, or bending while actively engaging in printed documents.

As indicated by prior studies [Sellen 2003, Steimle 2012a], knowledge work requires concurrent interaction with physical as well as digital objects. However, in contrast to the physical desktop workspace, interaction with the digital world is limited to the indirect and time-multiplex [Fitzmaurice 1995] manipulation of objects through mice and keyboards. Therefore, the overarching goal of this thesis was to extend the digital computation and interaction to the physical workspace of knowledge workers, towards more efficient, intuitive, and direct manipulation of digital objects. To this end, the present thesis explored novel interaction concepts that exploit flexible manipulation of physical objects. The concepts are designed in response to the challenges emerging from digitalization of two salient and key components of the knowledge work ensemble – namely, table and paper.

In this concluding chapter, a summary of the contributions of presented in this thesis is given. Then, pointers to future research are provided that may lead in several directions.

5.1 Summary

This thesis followed two main research directions, coined as table- and paper-centric, that we briefly recall the main outcomes of each in the following.

5.1.1 Table-centric Approach

Initial User Study

The first direction addressed challenges stemming from the integration of interactive tabletops into the desktop workspaces. More precisely, with our contributions in this direction, we aimed to inform the design of tabletop systems that allow a cohabitation of both physical and digital objects on one unique surface. Following a user-centered design approach, we first empirically investigated how people concurrently interact with a set of physical and digital documents on digital tabletops. We analyzed hybrid physical-digital usage and behavioral patterns while grouping, searching, and browsing documents. We also analyzed how users deal with occlusion of screen contents by physical documents. More explicitly, we took a close look at how occlusion influences the activities of selecting, accessing, zooming, moving, and grouping digital items.

Among others, the study showed that physical documents were mainly used above the tabletop surface and interaction with them was found to be easy, intuitive, and less mentally demanding. On the contrary, interaction with digital objects and piles was found to be visually demanding and was performed single-handedly most of the time. The study also revealed that people organize both types of documents in the form of a hybrid pile – i.e., a spatially integrated representation of physical and digital documents that are semantically related. In addition, we found that users are willing to physically occlude digital contents to better manage their workspace and make meaningful collections. However, it was cumbersome, when users had to access occluded digital contents. When coping with occlusion, we observed that users had effective strategies that rely on bimanual interaction.

Findings of this study were compiled into two sets of requirements – one related to hybrid piling issues and the other related to the occlusion issues – that provided a solid grounding for the design of two tabletop interfaces: *PaperTop* and *ObjecTop*.

PaperTop

We summarized the findings of the study into several requirements that provided foundations for the design of the PaperTop interface – i.e., a user interface for hybrid media piling on interactive tabletops. It consists of a set of interaction and visualization concepts to facilitate manipulation of hybrid piles. Two sets of such concepts can be distinguished. First, the PaperTop interface continuously provides unambiguous feedback about the hybrid grouping of documents using a soap bubble-like visualization. The soup bubble is a lightweight and unobtrusive visualization that flexibly adapts its shape to any arbitrary spatial arrangement of documents in a hybrid pile. Second, by leveraging the tangibility of physical documents constituting a hybrid pile, users are able to manipulate the entire pile in a coherent manner. This means that once the user wants to obtain an overview of the entire pile, she can easily arrange physical documents in various representational forms, then digital documents contained in the pile will automatically rearrange, accordingly.

The PaperTop interface was qualitatively evaluated in an early user feedback session with five HCI experts. Results showed that PaperTop is easy to understand and all experts were able to quickly pick up its interaction style. The spatial proximity that was used as the underlying principle to infer an aggregation of documents in group was found to be intuitive and provided an implicit way of defining hybrid groups. The experts' feedback highlighted that the flexible reorganization scheme of the PaperTop interface enabled fluid transition between neat hybrid piles and full juxtaposition, by displaying the documents linearly in a vertical or horizontal matter, or by fanning them out. Based on findings, we concluded that the PaperTop interface successfully exploited the flexibility and fluidity of interaction with physical piles to hybrid ones, and therefore, offered a rich user experience.

ObjecTop

Based on an extensive set of requirements, we designed ObjecTop, an occlusion management framework that supports awareness, access, and organization of occluded objects. Its design takes both inconvenient and desirable aspects of occlusion into accounts. In particular, it reflects intentional occlusion of digital items properly. At the same time, it supports access to occluded items by visualizing an interactive proxy-like representation for each occluded item placed on the nearest edge of the occluder. Users then are able to access occluded items by dragging out the respective proxy. While dragging out, the proxy visualized various levels of detail about the occluded item before fully accessing the actual item. This provides a gradual sneak-peek into the occluded item before resolving the occlusion. In addition, we contributed a pressure-based technique – namely, PressView – that leverages physical manipulation of physical objects. Using this technique, users are able to press down the physical occluder lightly such that the proxies would temporary squeeze out – i.e., showing more levels of details.

The ObjecTop framework also provided a collection of bimanual gestures that facilitates interaction with digital objects across physical occluders. This collection included a technique for moving digital items across physical occluders, zooming digital items remotely on any empty area on the tabletop surface, hybrid piling or hiding digital items under physical ones, and persistently binding digital items to physical ones.

Through an iterative design process, we evaluated and refined the design of ObjecTop. We first examined its main concepts in an exploratory user study. It was found that for complex physical occluders – such as books or laptops –dragging out proxies was mostly used to resolve occlusion. Regarding various levels of detail provided by proxies, we observed that the icon representation that consumes less space and introduces less visual clutter was often considered most suitable for conveying awareness. The results of this study left us with several efficiency-related questions that were formulated as hypothesis to be evaluated in a controlled experiment. More explicitly, we quantitatively evaluated the ObjecTop techniques in terms of efficiency, effectiveness and user experience compared to an unaided tabletop system that users could only move or lift physical objects to resolve occlusion.

The findings from the controlled experiment revealed that ObjecTop performed significantly faster than the unaided tabletop system for the searching task used as benchmark. Moreover, it significantly reduced the number of interactions with physical objects. This was because ObjecTop enabled users to employ an efficient combination of strategies based on the physical properties of the occluder. E.g, users used the drag-out technique to resolve occlusion caused by heavy and hardly movable occluders such as laptops. As to lightweight physical occluders – such as notepads, paper sheets, etc. – users easily moved or lifted them to access underlying items. This efficient combination of two strategies resulted in less physically demanding interaction with objects, lower perceived effort and frustration, while manipulating digital items.

We believe that both ObjecTop and PaperTop interfaces addressed several important issues stemming from integrating tabletops into desktop settings. Contributions proposed in this research direction can serve as a first step toward making interactive tabletops *real tables* that are seamlessly integrated into our everyday life.

5.1.2 Paper-centric Approach

In this research direction, we approached the knowledge work from a paper-centric perspective. We were particularly motivated by the concept of ePaper displays [Co 2008] that allows a paper-like manipulation of digital contents displayed on a flexible display. More precisely, this research direction aimed at investigating how people would interact with digital contents displayed on ultra-thin flexible displays that can be folded and rolled, as well as how natural practices of working with paper such as folding, flipping, and rolling can be used as new modes of interaction.

Resizable Display Concepts

We proposed two novel display concepts – namely FoldMe, a dual-sided foldable display concept, and Xpaaand, a rollable display concept – that enable shape and size to be modulated through folding and rolling techniques. We systematically explored the design space of both concepts. The FoldMe design space covered various forms, types, and numbers of folding. In the design space of Xpaaand, in addition to left, right, and symmetrical rolling types, we defined two metaphors (locked-in-hand and locked-in-viewport) so that digital contents can be displayed while (un)rolling the display. Since it is very likely that future flexible displays will be multitouch-enabled, we considered various combinations of touch-and-fold and touch-and-roll in the design spaces.

Interaction Principles and Techniques

Based on the design spaces, we contributed a set of interaction concepts that leveraged resizing and reshaping of the display as an input modality. These concepts were designed to facilitate the accomplishment of a set of basic tasks – such as browsing, zooming, navigating in hierarchies, application switching, and multitasking. The concepts were implemented in the form of fully functional prototypes using a passive display approach: passive surfaces (e.g., plastic, paper, or cardboard) are virtually transformed into touch-enabled flexible displays with the help of interactive distortion-free top-projection and precise (optical) tracking of these surfaces and of the users' fingers. This enabled us to test the effectiveness of the interaction concepts in two user studies. We particularly focused on obtaining qualitative insights on how people apprehend physical resizing of displays as a novel input modality.

Evaluations

We evaluated each display concept separately. Both user studies showed that physical resizing of displays on-the-fly can offer a delightful user experience. In FoldMe, folding was perceived to be very natural and quick, in particular compared to flipping the whole device to access the display located on the back side. Dynamic and continuous resizing of the display as afforded by Xpaaand was found efficient and intuitive in various tasks such as searching and comparing items in a list. In general, enlarging the display was interpreted as *offering options* or *collating the viewport*, depending on the task at hand. Similarly, shrinking the display resembled either *committing changes* or *consolidating the viewport*. A common finding in both user studies was that users did not prefer to use physical resizing (be it through rolling or folding) to accomplish tasks that required frequent collapsing and enlarging of the display. There was a strong consensus that for repetitive actions or frequent application switching, users preferred to use input modalities – such as direct touch input – that were physically less demanding. In summary, the user studies confirmed that physical resizing of the display has a great potential to enhance the interaction with digital contents.

5.2 Potential Directions of Future Work

In the following, we first sketch potential ways that the two research directions explored in this thesis can be integrated. We moreover identify two broader research directions that are discussed next.

Integrating Interactive Tabletops and Resizable Displays

As it was mentioned in the introduction, tables and paper-based media play a crucial role in knowledge work activities. While the contributions presented in this thesis show how interactive tabletops and resizable paper-like displays can be effectively integrated into the desktop workspaces *separately*, future work should further investigate the interplay of table-centric and paper-centric approaches in supporting of knowledge workers. We envisage two concrete ways that the tight integration of both research directions can facilitate knowledge work activities.

The first way is to support knowledge workers in activities that take place in the physical desktop setting. In these activities, foldable and rollable displays can complement other ways of displaying and interacting with digital contents on tabletop displays (cf. table-centric concepts). As a concrete example, active reading is a critical task of knowledge workers [Adler 1998] that can benefit from such a multi-display desktop environment. As indicated by many prior studies, such activity is characterized by frequent transition between reading as the primary task, and a set of sub-tasks – such as annotation [Marshall 1997], content browsing [O'Hara 1997], and cross-referencing [Steimle 2009a]. Knowledge workers could use a foldable display for the main reading task, and allocate the sub-tasks and corresponding materials to other displays – the display located on the back side of the foldable display, tabletop screen surface, etc. – based on for instance frequency or ease of use.

In this way, knowledge workers would also be able to continue reading in different postural configurations [Hong 2012]: holding the foldable display with both hands in mid-air to read in a lean-back posture while supplementary materials are shown on the desktop monitor, or switch to a lean-over posture by placing the foldable display on the tabletop surface, where the user disposes of other digital and physical materials. This multi-display desktop workspace could potentially offer a high level of interconnectivity among reading-related tasks, and thus a richer user experience for knowledge workers than conventional settings. Future work would be needed to investigate how users interact in this augmented ensemble, and how digital media is used on resizable displays, tabletops, and the computer monitor. Important and in-depth research would be needed to find appropriate techniques and interaction concepts for seamless interplay of tabletop and multiple flexible displays. This concerns, e.g. the mapping between physical arrangements of devices and digital arrangements of documents, carefully balancing automatic and user-controlled actions, concepts for adapting content placement and presentation to the user's device arrangement, posture, focus device, etc.

The second domain in which a careful combination of table- and paper-centric contributions of this thesis can assist knowledge work activities is mobile work. As indicated in prior studies [Sellen 2003, Steimle 2012a], *mobile use* is one of the key affordances of paper. This is mainly true because paper is thin, lightweight and flexible so that it can be used in many mobile situations and physical places. The resizable displays (or ultimately mobile devices featuring such displays) can potentially act as *handhled information containers* that knowledge workers can use for carrying digital contents from their desktop to other physical places and the other way around. Such scenarios typically consist of three steps: collecting and copying information from the desktop workspace to a handheld flexible display, carrying it physically to another place, and finally, spreading information from the display to another physical desktop workspace.

While we investigated how knowledge workers can effectively interact with digital content on resizable displays on-the-go (i.e, the second step), the two other steps need further exploration. More explicitly, future work is needed to address how knowledge workers can effectively collect and transfer digital contents represented on a hybrid tabletop surface to a resizable handheld device and the other way around.

A "metamorphosis" rather than an interplay of paper-centric and table-centric approaches is inspired by a recent study by Steimle and Olberding [Steimle 2012c] that showed future rollable displays can act as a relatively large interactive surface for mobile scenarios. For instance, large rollout displays were found suitable for short on-the-fly meetings in a mobile setting [Steimle 2012c]. Such displays could ultimately take the (rolled-out) size of tabletops. Therefore, another direction of future work could consider how hybrid tabletop concepts presented in sections 2.3 and 3.3 can be tailored to situations, where the tabletop surface itself is a large handheld rollable display.

In-Situ Resolving of Occlusion Using Top-Projection

The occlusion management framework presented in chapter 3 aimed to resolve physical occlusion by providing additional means (e.g., proxies and halos) for accessing occluded items. In the evaluations, we have shown that this approach is practical. However, it introduces a certain amount of visual clutter produced by visualizing the proxies and halos. While the amount of such visual clutter turned out to be reasonable for mildly populated tabletop settings, it was distracting and space-consuming when many physical objects occluded the screen area. Another limitations leis in the fact that proxies (as iconized representation of occluded objects) only convey object type and name to users initially. Other basic information about occluded objects – such as size, shape, layout, and texture – provide important clues to find them that are not supported by the proxy concept.

To overcome these limitations, the display and interaction space could be extended to the surface of physical objects as occluders [Liao 2010], tabletop rims [Cheng 2010], and even the hands and forearms of users [Adachi 2013b]. Given the recent advances in projection and tracking technologies [Huber 2012], future hybrid tabletop systems can be equipped with a highly capable projector-camera unit (e.g., in the form of a desktop lamp [Linder 2010, Riemann 2013]) to display digital contents and to sense user input on the surface of physical objects. In summary, a promising direction of research is the use of (camera assisted) top projection for resolving occlusion *in situ*. One obvious yet unexplored research question to be addressed in this context is how an occluded digital item can be appropriately represented on the physical occluder.

Future of Interactive Surfaces: Reshapable and Stretchable Displays

Since the advent of the very first flexible ePaper in 1974, flexible display technology has undergone a profound revolution. Conceptual and technological advances in the late 2000s – such as Raedius rollable cell phone prototype in 2007, Samsung's FOLED flexible display in 2007, Nokia Morph and Kinetic concepts in 2008, and Sony's rollable display prototypes in 2010 – have sparked a new surge of interest in implementing this technology in mobile devices. These developments led to the recent research prototypes based on actual flexible displays (e.g., PaperPhone [Lahey 2011]) as well as commercial availability of mobile phones that feature curved flexible display (e.g., LG G Flex [1]) Therefore, handheld devices featuring bendable, rollable, or foldable devices that looked more like science fiction at the time of starting this thesis, are now becoming market-ready products and are expected to achieve mass adoption in the foreseeable future.

Based on the ever growing trend of developing flexible display materials and printed electronics [Crawford 2005], we believe that in the far future, displays will become even more paper-like so that they can be arbitrarily reshaped and customized. This means that future displays can be not only deformed by folding or rolling them along one dimension, but tailored to various custom shapes. This opens up even greater interaction possibilities compared to flexible displays with rigid substrate and expensive hardware components. Future displays may become so thin and malleable that they can be wrapped around non-planar objects and, potentially, embrace everyday objects around us [Vertegaal 2011]. The field of *arbitrarily reshapable displays* is an interesting direction that can set the stage for future research.

Corresponding to reshapable output, the issue of reshapable input surfaces must also be addressed. Work is emerging that aim to provide customizable flexible sensing for non-planar surfaces [Olberding 2013, Holman 2011]. At the time of this writing, a considerable number of important user- and technology-related challenges remain unexplored.

Roughly the same timespan that applies to reshapable displays also applies to *stretchable displays* (modulated by the uncertainties of technology evolution). They have come into further reach due to recent advances in combining elastic interconnects with discrete rigid organic or inorganic light-emitting diodes (OLEDs or LEDs) [Liang 2013] have led to create displays that are *stretchable*. We believe that such displays can be used in a wide variety of applications. One potential application domain that can largely benefit from such stretchable displays is wearable computing. Future work in this direction can explore for instance the concept of *e-textile*

[1] http://www.lg.com/us/mobile-phones/gflex

displays or *e-skin patch displays* as novel forms of wearable computers.

Both reshapable and stretchable displays seem to be attractive as part of everyday or special-purpose physical objects such as appliances (reshapble) and clothing or medical coating (stretchable). However, their affordances may rather quickly turn them into ubiquitous commodities. At this point at the latest, they will also become attractive in the context of knowledge work and hence as future research along the line of this thesis.

Appendix A

Appendix

A.1 NASA TLX Questionnaire

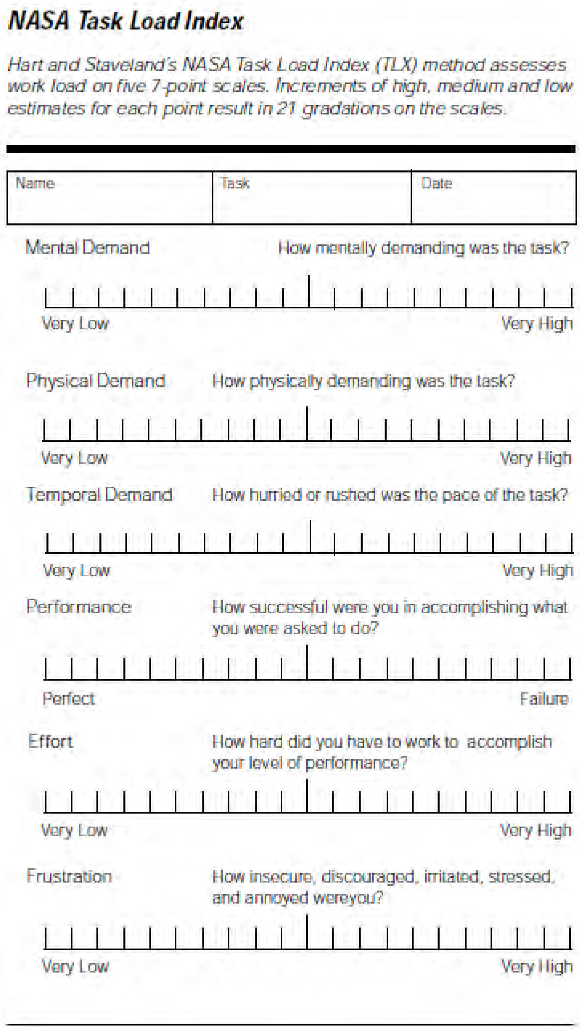

NASA Task Load Index

Hart and Staveland's NASA Task Load Index (TLX) method assesses work load on five 7-point scales. Increments of high, medium and low estimates for each point result in 21 gradations on the scales.

Name	Task	Date

Mental Demand — How mentally demanding was the task?

Very Low — Very High

Physical Demand — How physically demanding was the task?

Very Low — Very High

Temporal Demand — How hurried or rushed was the pace of the task?

Very Low — Very High

Performance — How successful were you in accomplishing what you were asked to do?

Perfect — Failure

Effort — How hard did you have to work to accomplish your level of performance?

Very Low — Very High

Frustration — How insecure, discouraged, irritated, stressed, and annoyed wereyou?

Very Low — Very High

Figure A.1: NASA TLX Questionnaire

Bibliography

[Adachi 2013a] Takamasa Adachi, Seiya Koura, Fumihisa Shibata and Asako Kimura. *Forearm Menu: Using Forearm As Menu Widget on Tabletop System.* In Proceedings of the 2013 ACM International Conference on Interactive Tabletops and Surfaces, ITS '13, pages 333–336, New York, NY, USA, 2013. ACM. (Cited on page 59.)

[Adachi 2013b] Takamasa Adachi, Seiya Koura, Fumihisa Shibata and Asako Kimura. *Forearm Menu: Using Forearm As Menu Widget on Tabletop System.* In Proceedings of the 2013 ACM International Conference on Interactive Tabletops and Surfaces, ITS '13, pages 333–336, New York, NY, USA, 2013. ACM. (Cited on page 172.)

[Adler 1998] Annette Adler, Anuj Gujar, Beverly L. Harrison, Kenton O'Hara and Abigail Sellen. *A Diary Study of Work-related Reading: Design Implications for Digital Reading Devices.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '98, pages 241–248, New York, NY, USA, 1998. ACM Press/Addison-Wesley Publishing Co. (Cited on page 170.)

[Agarawala 2006] Anand Agarawala and Ravin Balakrishnan. *Keepin' It Real: Pushing the Desktop Metaphor with Physics, Piles and the Pen.* In Proceedings of the 24th SIGCHI Conference on Human Factors in Computing Systems, CHI '06, pages 1283–1292, New York, NY, USA, 2006. ACM. (Cited on pages 19 and 20.)

[Aitenbichler 2007] Erwin Aitenbichler, Jussi Kangasharju and Max Mühlhäuser. *MundoCore: A Light-weight Infrastructure for Pervasive Computing.* Pervasive Mob. Comput., vol. 3, no. 4, pages 332–361, August 2007. (Cited on page 47.)

[Balakrishnan 1999] Ravin Balakrishnan, George Fitzmaurice, Gordon Kurtenbach and Karan Singh. *Exploring Interactive Curve and Surface Manipulation Using a Bend and Twist Sensitive Input Strip.* In Proceedings of the 1999 Symposium on Interactive 3D Graphics, I3D '99, pages 111–118, New York, NY, USA, 1999. ACM. (Cited on page 119.)

[Baudisch 2004] Patrick Baudisch and Carl Gutwin. *Multiblending: Displaying Overlapping Windows Simultaneously Without the Drawbacks of Alpha Blending.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '04, pages 367–374, New York, NY, USA, 2004. ACM. (Cited on pages 56 and 57.)

[Bauer 2005] Dan Bauer, Pierre Fastrez and Jim Hollan. *Spatial Tools for Managing Personal Information Collections.* In Proceedings of the 38th Hawaii International Conference on System Sciences, 2005. (Cited on page 19.)

[Bederson 1996] Benjamin B. Bederson, James D. Hollan, Ken Perlin, Jonathan Meyer, David Bacon and George Furnas. *Pad++: A Zoomable Graphical Sketchpad For Exploring Alternate Interface Physics.* Journal of Visual Languages & Computing, vol. 7, no. 1, pages 3 – 32, 1996. (Cited on pages 19, 114 and 136.)

[Bell 2001] Blaine Bell, Steven Feiner and Tobias Höllerer. *View Management for Virtual and Augmented Reality.* In Proceedings of the 14th Annual ACM Symposium on User Interface Software and Technology, UIST '01, pages 101–110, New York, NY, USA, 2001. ACM. (Cited on page 56.)

[Blinn 1982] James F. Blinn. *A Generalization of Algebraic Surface Drawing.* ACM Trans. Graph., vol. 1, no. 3, pages 235–256, July 1982. (Cited on page 49.)

[Bondarenko 2005] Olha Bondarenko and Ruud Janssen. *Documents at Hand: Learning from Paper to Improve Digital Technologies.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '05, pages 121–130, New York, NY, USA, 2005. ACM. (Cited on page 6.)

[Brandl 2009] Peter Brandl, Jakob Leitner, Thomas Seifried, Michael Haller, Bernard Doray and Paul To. *Occlusion-aware Menu Design for Digital Tabletops.* In CHI '09 Extended Abstracts on Human Factors in Computing Systems, CHI EA '09, pages 3223–3228, New York, NY, USA, 2009. ACM. (Cited on page 58.)

[Burstyn 2013] Jesse Burstyn, Amartya Banerjee and Roel Vertegaal. *FlexView: An Evaluation of Depth Navigation on Deformable Mobile Devices.* In Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction, TEI '13, pages 193–200, New York, NY, USA, 2013. ACM. (Cited on pages 124, 138 and 140.)

[Chen 2008] Nicholas Chen, Francois Guimbretiere, Morgan Dixon, Cassandra Lewis and Maneesh Agrawala. *Navigation Techniques for Dual-display e-Book Readers*. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '08, pages 1779–1788, New York, NY, USA, 2008. ACM. (Cited on pages 118, 126 and 160.)

[Cheng 2010] Kelvin Cheng, Christian Müller-Tomfelde and Natalie Ruiz. *Interactive Tabletops with Non-interactive Rims*. In Proceedings of the 22Nd Conference of the Computer-Human Interaction Special Interest Group of Australia on Computer-Human Interaction, OZCHI '10, pages 296–299, New York, NY, USA, 2010. ACM. (Cited on page 172.)

[Co 2008] Elise Co and Nikita Pashenkov. *Emerging Display Technologies for Organic User Interfaces*. Commun. ACM, vol. 51, no. 6, pages 45–47, June 2008. (Cited on pages 118 and 169.)

[Cockburn 2009] Andy Cockburn, Amy Karlson and Benjamin B. Bederson. *A Review of Overview+Detail, Zooming, and Focus+Context Interfaces*. ACM Comput. Surv., vol. 41, no. 1, pages 2:1–2:31, January 2009. (Cited on page 136.)

[Cotting 2006] Daniel Cotting and Markus Gross. *Interactive Environment-aware Display Bubbles*. In Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology, UIST '06, pages 245–254, New York, NY, USA, 2006. ACM. (Cited on page 61.)

[Crawford 2005] G. Crawford. Flexible flat panel displays. Wiley Series in Display Technology. Wiley, 2005. (Cited on pages 9, 119 and 173.)

[Davidson 2008] Philip L. Davidson and Jefferson Y. Han. *Extending 2D Object Arrangement with Pressure-sensitive Layering Cues*. In Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology, UIST '08, pages 87–90, New York, NY, USA, 2008. ACM. (Cited on page 89.)

[Dijkstra 2011] Rob Dijkstra, Christopher Perez and Roel Vertegaal. *Evaluating Effects of Structural Holds on Pointing and Dragging Performance with Flexible Displays*. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '11, pages 1293–1302, New York, NY, USA, 2011. ACM. (Cited on page 124.)

[Everitt 2008] 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. IEEE, 2008. (Cited on pages 23, 25 and 26.)

[Fishkin 2000] Kenneth P. Fishkin, Anuj Gujar, Beverly L. Harrison, Thomas P. Moran and Roy Want. *Embodied User Interfaces for Really Direct Manipulation*. Commun. ACM, vol. 43, no. 9, pages 74–80, September 2000. (Cited on page 119.)

[Fitts 1992] Paul M. Fitts. *The information capacity of the human motor system in controlling the amplitude of movement*. Experimental Psychology, vol. 47, pages 262–269, 1992. (Cited on page 127.)

[Fitzmaurice 1995] George W. Fitzmaurice, Hiroshi Ishii and William A. S. Buxton. *Bricks: Laying the Foundations for Graspable User Interfaces*. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '95, pages 442–449, New York, NY, USA, 1995. ACM Press/Addison-Wesley Publishing Co. (Cited on pages 1, 2 and 165.)

[Freeman 2013] Euan Freeman and Stephen Brewster. *Messy Tabletops: Clearing Up the Occlusion Problem*. In CHI '13 Extended Abstracts on Human Factors in Computing Systems, CHI EA '13, pages 1515–1520, New York, NY, USA, 2013. ACM. (Cited on page 62.)

[Furumi 2012] Genki Furumi, Daisuke Sakamoto and Takeo Igarashi. *SnapRail: A Tabletop User Interface Widget for Addressing Occlusion by Physical Objects*. In Proceedings of the 2012 ACM International Conference on Interactive Tabletops and Surfaces, ITS '12, pages 193–196, New York, NY, USA, 2012. ACM. (Cited on pages 60 and 102.)

[Gallant 2008] David T. Gallant, Andrew G. Seniuk and Roel Vertegaal. *Towards More Paper-like Input: Flexible Input Devices for Foldable Interaction Styles*. In Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology, UIST '08, pages 283–286, New York, NY, USA, 2008. ACM. (Cited on page 120.)

[Girouard 2012] Audrey Girouard, Aneesh Tarun and Roel Vertegaal. *DisplayStacks: Interaction Techniques for Stacks of Flexible Thin-film Displays*.

In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '12, pages 2431–2440, New York, NY, USA, 2012. ACM. (Cited on page 123.)

[Guiard 1987] Yves Guiard. *Asymmetric Division of Labor in Human Skilled Bimanual Action: The Kinematic Chain as a Model.* Journal of Motor Behavior, vol. 19, pages 486–517, 1987. (Cited on pages 30 and 81.)

[Haller 2006] Michael Haller, Daniel Leithinger, Jakob Leitner, Thomas Seifried, Peter Brandl, Jürgen Zauner and Mark Billinghurst. *The Shared Design Space.* In ACM SIGGRAPH 2006 Emerging Technologies, SIGGRAPH '06, New York, NY, USA, 2006. ACM. (Cited on pages 22, 23 and 24.)

[Hancock 2009] Mark Hancock, Otmar Hilliges, Christopher Collins, Dominikus Baur and Sheelagh Carpendale. *Exploring Tangible and Direct Touch Interfaces for Manipulating 2D and 3D Information on a Digital Table.* In Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces, ITS '09, pages 77–84, New York, NY, USA, 2009. ACM. (Cited on page 22.)

[Harrison 1995] Beverly L. Harrison, Gordon Kurtenbach and Kim J. Vicente. *An Experimental Evaluation of Transparent User Interface Tools and Information Content.* In Proceedings of the 8th Annual ACM Symposium on User Interface and Software Technology, UIST '95, pages 81–90, New York, NY, USA, 1995. ACM. (Cited on page 56.)

[Harrison 1998] Beverly L. Harrison, Kenneth P. Fishkin, Anuj Gujar, Carlos Mochon and Roy Want. *Squeeze Me, Hold Me, Tilt Me! An Exploration of Manipulative User Interfaces.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '98, pages 17–24, New York, NY, USA, 1998. ACM Press/Addison-Wesley Publishing Co. (Cited on page 119.)

[Hart 2006] Sandra G. Hart. *NASA-TASK LOAD INDEX (NASA-TLX); 20 YEARS LATER.* In Human Factors and Ergonomics Society 50th Annual Meeting, pages 904–908, 2006. (Cited on page 105.)

[Hartmann 2010] Björn Hartmann, Meredith Ringel Morris, Hrvoje Benko and Andrew D. Wilson. *Pictionaire: Supporting Collaborative Design Work by Integrating Physical and Digital Artifacts.* In Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work, CSCW '10, pages

421–424, New York, NY, USA, 2010. ACM. (Cited on pages 22, 23, 24 and 38.)

[Herkenrath 2008] Gero Herkenrath, Thorsten Karrer and Jan Borchers. *Twend: Twisting and Bending As New Interaction Gesture in Mobile Devices.* In CHI '08 Extended Abstracts on Human Factors in Computing Systems, CHI EA '08, pages 3819–3824, New York, NY, USA, 2008. ACM. (Cited on page 120.)

[Hinckley 2003] Ken Hinckley. *Synchronous Gestures for Multiple Persons and Computers.* In Proceedings of the 16th Annual ACM Symposium on User Interface Software and Technology, UIST '03, pages 149–158, New York, NY, USA, 2003. ACM. (Cited on page 126.)

[Hinckley 2009] Ken Hinckley, Morgan Dixon, Raman Sarin, Francois Guimbretiere and Ravin Balakrishnan. *Codex: A Dual Screen Tablet Computer.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '09, pages 1933–1942, New York, NY, USA, 2009. ACM. (Cited on pages 126 and 160.)

[Holman 2005] David Holman, Roel Vertegaal, Mark Altosaar, Nikolaus Troje and Derek Johns. *Paper Windows: Interaction Techniques for Digital Paper.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '05, pages 591–599, New York, NY, USA, 2005. ACM. (Cited on pages 123 and 149.)

[Holman 2008] David Holman and Roel Vertegaal. *Organic User Interfaces: Designing Computers in Any Way, Shape, or Form.* Commun. ACM, vol. 51, no. 6, pages 48–55, June 2008. (Cited on page 8.)

[Holman 2011] David Holman and Roel Vertegaal. *TactileTape: Low-cost Touch Sensing on Curved Surfaces.* In Proceedings of the 24th Annual ACM Symposium Adjunct on User Interface Software and Technology, UIST '11 Adjunct, pages 17–18, New York, NY, USA, 2011. ACM. (Cited on page 173.)

[Hong 2012] Matthew Hong, Anne Marie Piper, Nadir Weibel, Simon Olberding and James Hollan. *Microanalysis of Active Reading Behavior to Inform Design of Interactive Desktop Workspaces.* In Proceedings of the 2012 ACM International Conference on Interactive Tabletops and Surfaces, ITS '12, pages 215–224, New York, NY, USA, 2012. ACM. (Cited on page 171.)

[Hornecker 2006] Eva Hornecker and Jacob Buur. *Getting a Grip on Tangible Interaction: A Framework on Physical Space and Social Interaction.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '06, pages 437–446, New York, NY, USA, 2006. ACM. (Cited on pages 2 and 7.)

[Huber 2012] Jochen Huber, Jürgen Steimle, Chunyuan Liao, Qiong Liu and Max Mühlhäuser. *LightBeam: Interacting with Augmented Real-world Objects in Pico Projections.* In Proceedings of the 11th International Conference on Mobile and Ubiquitous Multimedia, MUM '12, pages 16:1–16:10, New York, NY, USA, 2012. ACM. (Cited on page 172.)

[Ishak 2004] Edward W. Ishak and Steven K. Feiner. *Interacting with Hidden Content Using Content-aware Free-space Transparency.* In Proceedings of the 17th Annual ACM Symposium on User Interface Software and Technology, UIST '04, pages 189–192, New York, NY, USA, 2004. ACM. (Cited on page 57.)

[Ishii 1997] Hiroshi Ishii and Brygg Ullmer. *Tangible Bits: Towards Seamless Interfaces Between People, Bits and Atoms.* In Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems, CHI '97, pages 234–241, New York, NY, USA, 1997. ACM. (Cited on pages 2, 7 and 38.)

[Jacob 2008] Robert J.K. Jacob, Audrey Girouard, Leanne M. Hirshfield, Michael S. Horn, Orit Shaer, Erin Treacy Solovey and Jamie Zigelbaum. *Reality-based Interaction: A Framework for post-WIMP Interfaces.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '08, pages 201–210, New York, NY, USA, 2008. ACM. (Cited on page 2.)

[Javed 2011] Waqas Javed, KyungTae Kim, Sohaib Ghani and Niklas Elmqvist. *Evaluating Physical/Virtual Occlusion Management Techniques for Horizontal Displays.* In Proceedings of the 13th IFIP TC 13 International Conference on Human-computer Interaction - Volume Part III, INTERACT'11, pages 391–408, Berlin, Heidelberg, 2011. Springer-Verlag. (Cited on pages 60, 69, 79, 97 and 102.)

[Jordà 2007] Sergi Jordà, Günter Geiger, Marcos Alonso and Martin Kaltenbrunner. *The reacTable: Exploring the Synergy Between Live Music Performance and Tabletop Tangible Interfaces.* In Proceedings of the 1st International

Conference on Tangible and Embedded Interaction, TEI '07, pages 139–146, New York, NY, USA, 2007. ACM. (Cited on pages 5, 22, 23, 38 and 87.)

[Khalilbeigi 2009] Mohammadreza Khalilbeigi, Jürgen Steimle and Max Mühlhäuser. *Interaction Support for Hybrid Groups of Paper and Digital Documents on Tabletops*. In ACM International Conference on Interactive Tabletops and Surfaces (Adjunct CD). ACM, 2009. (Cited on pages 10 and 15.)

[Khalilbeigi 2010] Mohammadreza Khalilbeigi, Jürgen Steimle and Max Mühlhäuser. *Interaction Techniques for Hybrid Piles of Documents on Interactive Tabletops*. In CHI '10 Extended Abstracts on Human Factors in Computing Systems, CHI EA '10, pages 3943–3948, New York, NY, USA, 2010. ACM. (Cited on pages 10, 15, 40, 41, 42, 43 and 44.)

[Khalilbeigi 2011] Mohammadreza Khalilbeigi, Roman Lissermann, Max Mühlhäuser and Jürgen Steimle. *Xpaaand: Interaction Techniques for Rollable Displays*. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '11, pages 2729–2732, New York, NY, USA, 2011. ACM. (Cited on pages 10 and 117.)

[Khalilbeigi 2012a] Mohammadreza Khalilbeigi, Roman Lissermann, Wolfgang Kleine and Jürgen Steimle. *FoldMe: Interacting with Double-sided Foldable Displays*. In Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction, TEI '12, pages 33–40, New York, NY, USA, 2012. ACM. (Cited on pages 10 and 117.)

[Khalilbeigi 2012b] Mohammadreza Khalilbeigi, Patrik Schmittat, Max Mühlhäuser and Jürgen Steimle. *Occlusion-aware Interaction Techniques for Tabletop Systems*. In CHI '12 Extended Abstracts on Human Factors in Computing Systems, CHI EA '12, pages 2531–2536, New York, NY, USA, 2012. ACM. (Cited on pages 10 and 55.)

[Khalilbeigi 2013] Mohammadreza Khalilbeigi, Jürgen Steimle, Jan Riemann, Niloofar Dezfuli, Max Mühlhäuser and James D. Hollan. *ObjecTop: Occlusion Awareness of Physical Objects on Interactive Tabletops*. In Proceedings of the 2013 ACM International Conference on Interactive Tabletops and Surfaces, ITS '13, pages 255–264, New York, NY, USA, 2013. ACM. (Cited on pages 10 and 55.)

[Kidd 1994] Alison Kidd. *The Marks Are on the Knowledge Worker.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '94, pages 186–191, New York, NY, USA, 1994. ACM. (Cited on page 3.)

[Kildal 2012a] Johan Kildal, Susanna Paasovaara and Viljakaisa Aaltonen. *Kinetic Device: Designing Interactions with a Deformable Mobile Interface.* In CHI '12 Extended Abstracts on Human Factors in Computing Systems, CHI EA '12, pages 1871–1876, New York, NY, USA, 2012. ACM. (Cited on page 121.)

[Kildal 2012b] Johan Kildal and Graham Wilson. *Feeling It: The Roles of Stiffness, Deformation Range and Feedback in the Control of Deformable Ui.* In Proceedings of the 14th ACM International Conference on Multimodal Interaction, ICMI '12, pages 393–400, New York, NY, USA, 2012. ACM. (Cited on page 121.)

[Kildal 2013] Johan Kildal, Andrés Lucero and Marion Boberg. *Twisting Touch: Combining Deformation and Touch As Input Within the Same Interaction Cycle on Handheld Devices.* In Proceedings of the 15th International Conference on Human-computer Interaction with Mobile Devices and Services, MobileHCI '13, pages 237–246, New York, NY, USA, 2013. ACM. (Cited on page 124.)

[Kirk 2009] David Kirk, Abigail Sellen, Stuart Taylor, Nicolas Villar and Shahram Izadi. *Putting the Physical into the Digital: Issues in Designing Hybrid Interactive Surfaces.* In Proceedings of the 23rd British HCI Group Annual Conference on People and Computers: Celebrating People and Technology, BCS-HCI '09, pages 35–44, Swinton, UK, UK, 2009. British Computer Society. (Cited on pages 5, 8, 22 and 38.)

[Kirsh 1995] David Kirsh. *The Intelligent Use of Space.* Artif. Intell., vol. 73, no. 1-2, pages 31–68, February 1995. (Cited on page 3.)

[Klemmer 2001] Scott R. Klemmer, Mark W. Newman, Ryan Farrell, Mark Bilezikjian and James A. Landay. *The Designers' Outpost: A Tangible Interface for Collaborative Web Site.* In Proceedings of the 14th Annual ACM Symposium on User Interface Software and Technology, UIST '01, pages 1–10, New York, NY, USA, 2001. ACM. (Cited on pages 23, 24 and 25.)

[Koffka 1935] Kurt Koffka. Principles of gestalt psychology. International library of psychology, philosophy, and scientific method. Harcourt, Brace and Company, 1935. (Cited on page 33.)

[Koike 2001] Hideki Koike, Yoichi Sato and Yoshinori Kobayashi. *Integrating Paper and Digital Information on EnhancedDesk: A Method for Realtime Finger Tracking on an Augmented Desk System*. ACM Transactions on Computer-Human Interaction, vol. 8, no. 4, pages 307–322, December 2001. (Cited on page 21.)

[Kruger 2003] Russell Kruger, Sheelagh Carpendale, Stacey D. Scott and Saul Greenberg. *How People Use Orientation on Tables: Comprehension, Coordination and Communication*. In Proceedings of the 2003 International ACM SIGGROUP Conference on Supporting Group Work, GROUP '03, pages 369–378, New York, NY, USA, 2003. ACM. (Cited on page 17.)

[Lahey 2011] Byron Lahey, Audrey Girouard, Winslow Burleson and Roel Vertegaal. *PaperPhone: Understanding the Use of Bend Gestures in Mobile Devices with Flexible Electronic Paper Displays*. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '11, pages 1303–1312, New York, NY, USA, 2011. ACM. (Cited on pages 122, 125, 149 and 173.)

[Lee 2008] Johnny C. Lee, Scott E. Hudson and Edward Tse. *Foldable Interactive Displays*. In Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology, UIST '08, pages 287–290, New York, NY, USA, 2008. ACM. (Cited on pages 123 and 149.)

[Lee 2010] Sang-Su Lee, Sohyun Kim, Bopil Jin, Eunji Choi, Boa Kim, Xu Jia, Daeeop Kim and Kun-pyo Lee. *How Users Manipulate Deformable Displays As Input Devices*. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '10, pages 1647–1656, New York, NY, USA, 2010. ACM. (Cited on pages 121, 122, 136 and 137.)

[Lee 2012] Sang-su Lee, Youn-kyung Lim and Kun-Pyo Lee. *Exploring the Effects of Size on Deformable User Interfaces*. In Proceedings of the 14th International Conference on Human-computer Interaction with Mobile Devices and Services Companion, MobileHCI '12, pages 89–94, New York, NY, USA, 2012. ACM. (Cited on page 121.)

[Leithinger 2007] Daniel Leithinger and Michael Haller. *Improving Menu Interaction for Cluttered Tabletop Setups with User-Drawn Path Menus*. In Tabletop, pages 121–128. IEEE Computer Society, 2007. (Cited on pages 61 and 62.)

[Leitner 2011] Jakob Leitner and Michael Haller. *Geckos: Combining Magnets and Pressure Images to Enable New Tangible-object Design and Interaction*. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '11, pages 2985–2994, New York, NY, USA, 2011. ACM. (Cited on page 89.)

[Liang 2013] Jiajie Liang, Lu Li, Xiaofan Niu, Zhibin Yu and Qibing Pei. *Elastomeric polymer light-emitting devices and displays*. Nature Photonics, vol. 7, pages 817–824, 2013. (Cited on page 173.)

[Liao 2010] Chunyuan Liao, Hao Tang, Qiong Liu, Patrick Chiu and Francine Chen. *FACT: Fine-grained Cross-media Interaction with Documents via a Portable Hybrid Paper-laptop Interface*. In Proceedings of the International Conference on Multimedia, MM '10, pages 361–370, New York, NY, USA, 2010. ACM. (Cited on pages 46 and 172.)

[Linder 2010] Natan Linder and Pattie Maes. *LuminAR: Portable Robotic Augmented Reality Interface Design and Prototype*. In Adjunct Proceedings of the 23Nd Annual ACM Symposium on User Interface Software and Technology, UIST '10, pages 395–396, New York, NY, USA, 2010. ACM. (Cited on page 172.)

[Lorensen 1987] William E. Lorensen and Harvey E. Cline. *Marching Cubes: A High Resolution 3D Surface Construction Algorithm*. In Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '87, pages 163–169, New York, NY, USA, 1987. ACM. (Cited on page 49.)

[Malone 1983] Thomas W. Malone. *How Do People Organize Their Desks?: Implications for the Design of Office Information Systems*. ACM Trans. Inf. Syst., vol. 1, no. 1, pages 99–112, January 1983. (Cited on pages 3, 18, 19, 39 and 165.)

[Mander 1992] Richard Mander, Gitta Salomon and Yin Yin Wong. *A &Ldquo;Pile&Rdquo; Metaphor for Supporting Casual Organization of Information*. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '92, pages 627–634, New York, NY, USA, 1992. ACM. (Cited on pages 18, 19, 39, 81, 83 and 165.)

[Marshall 1997] Catherine C. Marshall. *Annotation: From Paper Books to the Digital Library*. In Proceedings of the Second ACM International Conference on

Digital Libraries, DL '97, pages 131–140, New York, NY, USA, 1997. ACM. (Cited on page 171.)

[Marshall 2005a] Catherine C. Marshall and Sara Bly. *Turning the Page on Navigation.* In Proceedings of the 5th ACM/IEEE-CS Joint Conference on Digital Libraries, JCDL '05, pages 225–234, New York, NY, USA, 2005. ACM. (Cited on pages 3, 4 and 6.)

[Marshall 2005b] Catherine C. Marshall and Sara Bly. *Turning the Page on Navigation.* In Proceedings of the 5th ACM/IEEE-CS Joint Conference on Digital Libraries, JCDL '05, pages 225–234, New York, NY, USA, 2005. ACM. (Cited on page 118.)

[Mitsuhara 2010] Hiroyuki Mitsuhara, Yoneo Yano and Toshiyuki Moriyama. *Paper-top interface for supporting note-taking and its preliminary experiment.* In IEEE International Conference on Systems Man and Cybernetics (SMC), pages 3456–3462, 2010. (Cited on page 21.)

[Morris 2006] Meredith Ringel Morris, Anqi Huang, Andreas Paepcke and Terry Winograd. *Cooperative Gestures: Multi-user Gestural Interactions for Co-located Groupware.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '06, pages 1201–1210, New York, NY, USA, 2006. ACM. (Cited on page 112.)

[Muller-Tomfelde 2012] Christian Muller-Tomfelde and Morten Fjeld. *Tabletops: Interactive Horizontal Displays for Ubiquitous Computing.* Computer, vol. 45, no. 2, pages 78–81, Feb 2012. (Cited on page 5.)

[Norman 1986] Donald A. Norman and Stephen W. Draper. User centered system design; new perspectives on human-computer interaction. L. Erlbaum Associates Inc., Hillsdale, NJ, USA, 1986. (Cited on page 8.)

[O'Hara 1997] Kenton O'Hara and Abigail Sellen. *A Comparison of Reading Paper and On-line Documents.* In Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems, CHI '97, pages 335–342, New York, NY, USA, 1997. ACM. (Cited on pages 3, 6, 16 and 171.)

[Olberding 2013] Simon Olberding, Nan-Wei Gong, John Tiab, Joseph A. Paradiso and Jürgen Steimle. *A Cuttable Multi-touch Sensor.* In Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology, UIST '13, pages 245–254, New York, NY, USA, 2013. ACM. (Cited on page 173.)

[Ou 2013] Jifei Ou, Lining Yao, Daniel Tauber, Jürgen Steimle, Ryuma Niiyama and Hiroshi Ishii. *jamSheets: Thin Interfaces with Tunable Stiffness Enabled by Layer Jamming.* In Proceedings of the 8th International Conference on Tangible, Embedded and Embodied Interaction, TEI '14, pages 65–72, New York, NY, USA, 2013. ACM. (Cited on page 121.)

[Piper 2009] Anne Marie Piper and James D. Hollan. *Tabletop Displays for Small Group Study: Affordances of Paper and Digital Materials.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '09, pages 1227–1236, New York, NY, USA, 2009. ACM. (Cited on page 17.)

[Porter 1984] Thomas Porter and Tom Duff. *Compositing Digital Images.* In Proceedings of the 11th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '84, pages 253–259, New York, NY, USA, 1984. ACM. (Cited on page 56.)

[Ramos 2006] Gonzalo Ramos, George Robertson, Mary Czerwinski, Desney Tan, Patrick Baudisch, Ken Hinckley and Maneesh Agrawala. *Tumble! Splat! Helping Users Access and Manipulate Occluded Content in 2D Drawings.* In Proceedings of the Working Conference on Advanced Visual Interfaces, AVI '06, pages 428–435, New York, NY, USA, 2006. ACM. (Cited on page 57.)

[Rekimoto 1999] Jun Rekimoto and Masanori Saitoh. *Augmented Surfaces: A Spatially Continuous Work Space for Hybrid Computing Environments.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '99, pages 378–385, New York, NY, USA, 1999. ACM. (Cited on pages 22 and 24.)

[Riemann 2010] Jan Riemann. Multiprojektorsystem zur simulation hochauflösender papierühnlicher displays. Bachelor thesis, TU Darmstadt, 2010. (Cited on page 152.)

[Riemann 2013] Jan Riemann, Mohammadreza Khalilbeigi and Max Mühlhäuser. *PeriTop: Extending Back-projected Tabletops with Top-projected Peripheral Displays.* In Proceedings of the 2013 ACM International Conference on Interactive Tabletops and Surfaces, ITS '13, pages 349–352, New York, NY, USA, 2013. ACM. (Cited on page 172.)

[Robinson 1997] Peter Robinson, Dan Sheppard, Richard Watts, Robert Harding and Steven Lay. *Animated Paper Documents.* In Michael J. Smith, Gavriel Salvendy and Richard J. Koubek, editeurs, HCI (2), pages 655–658. Elsevier, 1997. (Cited on page 21.)

[Ryall 2004] Kathy Ryall, Clifton Forlines, Chia Shen and Meredith Ringel Morris. *Exploring the Effects of Group Size and Table Size on Interactions with Tabletop Shared-display Groupware.* In Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work, CSCW '04, pages 284–293, New York, NY, USA, 2004. ACM. (Cited on page 17.)

[Schwesig 2004] Carsten Schwesig, Ivan Poupyrev and Eijiro Mori. *Gummi: A Bendable Computer.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '04, pages 263–270, New York, NY, USA, 2004. ACM. (Cited on page 122.)

[Scott 2004] Stacey D. Scott, M. Sheelagh T. Carpendale and Kori M. Inkpen. *Territoriality in Collaborative Tabletop Workspaces.* In Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work, CSCW '04, pages 294–303, New York, NY, USA, 2004. ACM. (Cited on pages 17 and 29.)

[Sellen 1997] Abigail Sellen and Richard Harper. *Paper As an Analytic Resource for the Design of New Technologies.* In Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems, CHI '97, pages 319–326, New York, NY, USA, 1997. ACM. (Cited on pages 16 and 117.)

[Sellen 2003] Abigail J. Sellen and Richard H.R. Harper. The myth of the paperless office. MIT Press, Cambridge, MA, USA, 2003. (Cited on pages 4, 6, 14, 16, 25, 77, 117, 118, 165 and 171.)

[Shneiderman 1982] Ben Shneiderman. *The Future of Interactive Systems and Emergence of Direct Manipulation.* Behavior and Information Technology, vol. 1, pages 237–256, 1982. (Cited on page 1.)

[Song 2009] Hyunyoung Song, Tovi Grossman, George Fitzmaurice, François Guimbretiere, Azam Khan, Ramtin Attar and Gordon Kurtenbach. *PenLight: Combining a Mobile Projector and a Digital Pen for Dynamic Visual Overlay.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '09, pages 143–152, New York, NY, USA, 2009. ACM. (Cited on page 134.)

[Song 2011] Hyunyoung Song, Hrvoje Benko, Francois Guimbretiere, Shahram Izadi, Xiang Cao and Ken Hinckley. *Grips and Gestures on a Multi-touch Pen.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '11, pages 1323–1332, New York, NY, USA, 2011. ACM. (Cited on page 133.)

[Steimle 2009a] Jürgen Steimle. *Designing Pen-and-paper User Interfaces for Interaction with Documents.* In Proceedings of the 3rd International Conference on Tangible and Embedded Interaction, TEI '09, pages 197–204, New York, NY, USA, 2009. ACM. (Cited on page 171.)

[Steimle 2009b] Jürgen Steimle, Oliver Brdiczka and Max Mühlhäuser. *CoScribe: Integrating Paper and Digital Documents for Collaborative Knowledge Work.* IEEE Transactions on Learning Technologies, vol. 2, no. 3, pages 174–188, 2009. (Cited on pages 23, 24 and 25.)

[Steimle 2010a] Jürgen Steimle, Mohammadreza Khalilbeigi and Max Mühlhäuser. *Hybrid Groups of Printed and Digital Documents on Tabletops: A Study.* In CHI '10 Extended Abstracts on Human Factors in Computing Systems, CHI EA '10, pages 3271–3276, New York, NY, USA, 2010. ACM. (Cited on pages 10, 15 and 39.)

[Steimle 2010b] Jürgen Steimle, Mohammadreza Khalilbeigi, Max Mühlhäuser and James D. Hollan. *Physical and Digital Media Usage Patterns on Interactive Tabletop Surfaces.* In ACM International Conference on Interactive Tabletops and Surfaces, ITS '10, pages 167–176, New York, NY, USA, 2010. ACM. (Cited on pages 10, 15, 32, 33, 65 and 67.)

[Steimle 2012a] Jürgen Steimle. Pen-and-paper user interfaces. Springer, 2012. (Cited on pages 4, 6, 165 and 171.)

[Steimle 2012b] Jürgen Steimle, Roman Lissermann, Simon Olberding, Mohammadreza khalilbeigi, Wolfgang Kleine and Max Mühlhäuser. *Be-greifbare Interaktionen mit grössenveränderbaren Bildschirmen.* i-com Oldenbourg Verlag, vol. 11, no. 2, pages 26–31, 2012. (Cited on pages 10 and 129.)

[Steimle 2012c] Jürgen Steimle and Simon Olberding. *When Mobile Phones Expand into Handheld Tabletops.* In CHI '12 Extended Abstracts on Human Factors in Computing Systems, CHI EA '12, pages 271–280, New York, NY, USA, 2012. ACM. (Cited on page 172.)

[Steimle 2013] Jürgen Steimle, Andreas Jordt and Pattie Maes. *Flexpad: Highly Flexible Bending Interactions for Projected Handheld Displays.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '13, pages 237–246, New York, NY, USA, 2013. ACM. (Cited on pages 121, 123 and 149.)

[Sun 2008] Qibin Sun and Wolfgang Hürst. *Video Browsing on Handheld Devices—Interface Designs for the Next Generation of Mobile Video Players.* IEEE MultiMedia, vol. 15, no. 3, pages 76–83, July 2008. (Cited on page 141.)

[Tabard 2012] Aurélien Tabard, Juan David Hincapié Ramos and Jakob Bardram. *The eLabBench in the Wild: Supporting Exploration in a Molecular Biology Lab.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '12, pages 3051–3060, New York, NY, USA, 2012. ACM. (Cited on page 61.)

[Tabard 2013] Aurélien Tabard, Simon Gurn, Andreas Butz and Jakob Bardram. *A Case Study of Object and Occlusion Management on the eLabBench, a Mixed Physical/Digital Tabletop.* In Proceedings of the 2013 ACM International Conference on Interactive Tabletops and Surfaces, ITS '13, pages 251–254, New York, NY, USA, 2013. ACM. (Cited on page 61.)

[Tajika 2008] Taichi Tajika, Tomoko Yonezawa and Noriaki Mitsunaga. *Intuitive Page-turning Interface of e-Books on Flexible e-Paper Based on User Studies.* In Proceedings of the 16th ACM International Conference on Multimedia, MM '08, pages 793–796, New York, NY, USA, 2008. ACM. (Cited on page 122.)

[Tarun 2011] Aneesh P. Tarun, Byron Lahey, Audrey Girouard, Winslow Burleson and Roel Vertegaal. *Snaplet: Using Body Shape to Inform Function in Mobile Flexible Display Devices.* In CHI '11 Extended Abstracts on Human Factors in Computing Systems, CHI EA '11, pages 329–334, New York, NY, USA, 2011. ACM. (Cited on page 123.)

[Terrenghi 2007] Lucia Terrenghi, David Kirk, Abigail Sellen and Shahram Izadi. *Affordances for Manipulation of Physical Versus Digital Media on Interactive Surfaces.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '07, pages 1157–1166, New York, NY, USA, 2007. ACM. (Cited on pages 17 and 29.)

[Terrenghi 2008] Lucia Terrenghi, David Kirk, Hendrik Richter, Sebastian Krämer, Otmar Hilliges and Andreas Butz. *Physical Handles at the Interactive Surface: Exploring Tangibility and Its Benefits.* In Proceedings of the Working Conference on Advanced Visual Interfaces, AVI '08, pages 138–145, New York, NY, USA, 2008. ACM. (Cited on pages 22 and 23.)

[Toney 2006] Aaron Toney and Bruce H. Thomas. *Considering Reach in Tangible and Table Top Design.* In Proceedings of the First IEEE International Workshop on Horizontal Interactive Human-Computer Systems, TABLETOP '06, pages 57–58, Washington, DC, USA, 2006. IEEE Computer Society. (Cited on page 77.)

[Ullmer 1997] Brygg Ullmer and Hiroshi Ishii. *The metaDESK: Models and Prototypes for Tangible User Interfaces.* In Proceedings of the 10th Annual ACM Symposium on User Interface Software and Technology, UIST '97, pages 223–232, New York, NY, USA, 1997. ACM. (Cited on pages 5 and 22.)

[Underkoffler 1999] John Underkoffler and Hiroshi Ishii. *Urp: A Luminous-tangible Workbench for Urban Planning and Design.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '99, pages 386–393, New York, NY, USA, 1999. ACM. (Cited on page 5.)

[Vertegaal 2011] Roel Vertegaal. *The (Re)Usability of Everyday Computational Things: Why Industrial Design Will Be the New Interaction Design.* interactions, vol. 18, no. 1, pages 38–41, January 2011. (Cited on page 173.)

[Victor 2011] Bret Victor. *A Brief Rant on The Future of Interaction Design*, November 2011. (Cited on page 2.)

[Vogel 2007] Daniel Vogel and Patrick Baudisch. *Shift: A Technique for Operating Pen-based Interfaces Using Touch.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '07, pages 657–666, New York, NY, USA, 2007. ACM. (Cited on page 58.)

[Vogel 2009] Daniel Vogel, Matthew Cudmore, Géry Casiez, Ravin Balakrishnan and Liam Keliher. *Hand Occlusion with Tablet-sized Direct Pen Input.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '09, pages 557–566, New York, NY, USA, 2009. ACM. (Cited on pages 58 and 59.)

[Vogel 2010] Daniel Vogel and Ravin Balakrishnan. *Occlusion-aware Interfaces.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '10, pages 263–272, New York, NY, USA, 2010. ACM. (Cited on pages 58 and 59.)

[Vogel 2012] Daniel Vogel and Géry Casiez. *Hand Occlusion on a Multi-touch Tabletop.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '12, pages 2307–2316, New York, NY, USA, 2012. ACM. (Cited on pages 58 and 59.)

[Wagner 2007] Daniel Wagner and Dieter Schmalstieg. *ARToolkitPlus for Pose Tracking on Mobile Devices.* In Proceedings of Computer Vision Winter Workshop, pages 139–146, 2007. (Cited on page 47.)

[Waldner 2011] Manuela Waldner, Markus Steinberger, Raphael Grasset and Dieter Schmalstieg. *Importance-driven Compositing Window Management.* In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '11, pages 959–968, New York, NY, USA, 2011. ACM. (Cited on pages 56 and 57.)

[Watanabe 2007] Nayuko Watanabe, Motoi Washida and Takeo Igarashi. *Bubble clusters: an interface for manipulating spatial aggregation of graphical objects.* In Chia Shen, Robert J. K. Jacob and Ravin Balakrishnan, editeurs, UIST, pages 173–182. ACM, 2007. (Cited on pages 19, 20, 21, 26, 40 and 41.)

[Watanabe 2008] Jun-ichiro Watanabe, Arito Mochizuki and Youichi Horry. *Bookisheet: Bendable Device for Browsing Content Using the Metaphor of Leafing Through the Pages.* In Proceedings of the 10th International Conference on Ubiquitous Computing, UbiComp '08, pages 360–369, New York, NY, USA, 2008. ACM. (Cited on page 120.)

[Wellner 1991] Pierre Wellner. *The DigitalDesk Calculator: Tangible Manipulation on a Desk Top Display.* In Proceedings of the 4th Annual ACM Symposium on User Interface Software and Technology, UIST '91, pages 27–33, New York, NY, USA, 1991. ACM. (Cited on page 21.)

[Wellner 1993] Pierre Wellner. *Interacting with Paper on the DigitalDesk.* Commun. ACM, vol. 36, no. 7, pages 87–96, July 1993. (Cited on pages 5, 21, 22 and 26.)

[Whittaker 2001] Steve Whittaker and Julia Hirschberg. *The Character, Value, and Management of Personal Paper Archives.* ACM Trans. Comput.-Hum. Interact., vol. 8, no. 2, pages 150–170, June 2001. (Cited on pages 16, 18 and 19.)

[Wigdor 2006] Daniel Wigdor, Darren Leigh, Clifton Forlines, Samuel Shipman, John Barnwell, Ravin Balakrishnan and Chia Shen. *Under the Table Interaction.* In Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology, UIST '06, pages 259–268, New York, NY, USA, 2006. ACM. (Cited on page 58.)

[Zhang 2010] Zhengyou Zhang. *Estimating Projective Transformation Matrix (Collineation, Homography).* Rapport technique, Microsoft Research, 1993/2010. (Cited on page 150.)

Curriculum Vitae

since 06/2009	Technische Universität Darmstadt PhD candidate in the area of Human-Computer Interaction
2006 - 2009	RWTH-Aachen University, Aachen - Germany Field of Media Informatics - International Master Program
2001 - 2005	Kish University, Kish Island - Iran Bachelor of Science in Computer Engineering (software)
1997 - 2001	High School Kowsar, Tehran - Iran Mathematic-Physic diploma