DisplayStacks: Interaction Techniques for Stacks of Flexible Thin-Film Displays

Audrey Girouard[1,2], Aneesh P. Tarun[1] and Roel Vertegaal[1]

[1] Human Media Lab
School of Computing, Queen's University
Kingston, Ontario, K7L 3N6 Canada

[2] School of Information Technology
Carleton University
Ottawa, Ontario, K1S 5B6 Canada

audrey_girouard@carleton.ca, aneesh@cs.queensu.ca, roel@cs.queensu.ca

ABSTRACT
Stacking physical documents is one of the main forms of spatio-temporal organization of information. We present DisplayStacks, a system that enables physical stacking of digital documents via piles of flexible E Ink displays. With a conductive dot pattern sensor attached to the flexible display, we dynamically track the position and orientation of these displays in relation to one another. We introduce mechanisms for interacting with these physical stacks for access and manipulation of information using asymmetric bimanual interactions, such as providing contextual overviews. Initial user experiences indicate a preference for linear overlaps as a stacking configuration.

Author Keywords
Flexible Displays, E Ink, Stacks, Organic User Interfaces.

ACM Classification Keywords
H5.2 [Information interfaces and presentation]: User Interfaces.

General Terms
Design, Human Factors.

INTRODUCTION
Considerable effort has been put towards the development of paper computers that use multiple physical displays to represent windows into digital content. Wellner [43] early-on outlined some compelling reasons for this: 1) interaction techniques for digital documents are limited, and distinct from those used with paper; 2) the physicality of paper provides users with richer forms of interaction that are deeply embedded in their tactile-kinesthetic systems; 3) paper allows for efficient switching between multiple, parallel, documents; and 4) the reflective properties of paper provides for a superior reading experience. In the real world, paper documents are often stored in ways that provide distinct spatial correlates. For example, in stacks of physical documents, the lower the document, the older it is [29,30]. One advantage to the organization of information in physical stacks over the organization of digital windows on a computer desktop is that the location of windows remains perceptually stable. Furthermore, the spatial layout is not limited to a single small screen, but envelopes the user easing retrieval of documents pertinent to foreground or background tasks. While research in surface computing has explored software window stacking and hybrid physical-digital stacking techniques that replicate the ease of physical picking, shuffling, and stacking for document access and manipulation [38,41,44], bimanual interactions with virtual stacks of documents are problematic, as is selecting documents based on their virtual elevation [41].

CHI 2012, May 5-10, 2012, Austin, TX, USA.

Figure 1. Flexible displays allow for the physical manipulation of digital information, resembling paper.

With the advent of thin-film and flexible displays, which begin to approach the weight, thinness and flexibility of paper, it has become possible to reconsider some of the traditional physical organizational metaphors of paper in organizing digital documents (see Fig. 1). These new display technologies, such as Flexible E Ink and Flexible Organic Light Emitting Diodes (FOLEDs) [26], will potentially allow users to access and navigate digital information with methods that *physically* resemble paper documents, where this is beneficial. Flexible displays are particularly unique with respect to LCD displays in that they are very thin in the *z* dimension, making them particularly suitable to be organized in stacks. While prior research exists on physically grouping digital documents according to location (i.e. collocation) [17], there has been little to no work on physically stacking digital documents; current display technologies are simply too thick to be stacked in meaningful ways. So while flexibility is a feature of thin-film displays, we believe their extreme thinness and light weight are the key design elements that make flexible displays ideally suitable for stacked designs. If displays are too thick, stacks become too difficult to hold. When displays are heavy, it becomes difficult for mobile users to carry stacks of displays in their pocket or purse.

Sellen and Harper [36] point out that physical stacks and piles are often used in offices to organize and navigate documents on the fly. We therefore believe the use of stacks of thin-film displays may have distinct advantages over virtual windows for on-the-fly organization of documents: 1) Physical piles support crude ordering of related information while maintaining parallel access to multiple documents; 2) Stacked physical windows are not hidden and remain visible and tangible in the *z* dimension; and 3) Physical windows are more easily handled in groups, e.g., to serve as context-aware tool lenses into other windows in the stack.

Contribution

In this paper, we propose DisplayStacks, a system for physically organizing digital documents via stacks of thin-film E Ink displays. Our main contributions are: 1) the introduction of an electronic paper computer that uses multiple thin-film electronic paper displays, in which each computer window is represented by its own paper-like display and 2) techniques for interacting with stacks of such displays. While work exists on stacking digital displays, the light weight, thinness and flexibility of these displays allow for some novel interactions that mimic some benefits of paper document navigation. As these interactions require thin-film sensing technologies, we propose a method for dynamically tracking the position and orientation of displays relative to each other using a conductive dot pattern sensor layer affixed to the bottom of each flexible display. We report on an initial pilot study investigating tracking accuracy and user preferences.

RELATED WORK

Sellen and Harper [37] describe some characteristics of printed-paper that may explain its continued popularity. Rigid graphical user interfaces often feature input that is indirect, one-handed, and dependent on visual cues. By contrast, paper documents may: (1) be very thin, low-weight, allowing superior portability; (2) have many physical pages, each page pertaining only to a specific and physically delineated task context; (3) provide variable screen real estate that fits the current context of use; (4) use physical bend gestures with strong tactile and kinesthetic feedback for efficient navigation; (5) allow documents to be laid out, or collocated for easy access.

Stacking physical documents is one of the main forms of spatio-temporal organization of information. On physical desktops, piling is an advantageous method of organizing documents as it can be done on any flat surface. Piling is a lightweight, casual activity requiring little cognitive overhead [29]. It allows elements to be easily repositioned within a pile, or reorganized between piles [44]. Moreover, piles are most useful in tasks using visual features of documents [20].

Digital Piles

Mander et al. [30] introduced the pile metaphor to browse and manipulate sets of digital documents. Their 3D piles were presented either with a disheveled appearance to indicate user-created piles, or with a neat look, indicating system-created piles. Beaudouin-Lafon [5] created physically-inspired digital piles through Rotated Windows, displaying a top-view pile of “loose” documents. The user peeled back the top window to reveal the lower window. BumpTop [1] used pen-based interaction and visualization techniques to manipulate groups of electronic documents on a computer desktop using physically-simulated piles. Key elements of their prototype included the ability to toss and drag documents, create neat and messy piles, and to support pile browsing with a variety of widgets. Our work aims at creating a physical instantiation of these digital piles. With tabletop computing, Davidson and Han [8] extended the physical metaphor of layering to move between overlaid windows. Users could push the side of a virtual document to lift the other side, and move it atop.

In exploring differences between physical and digital media, Terrenghi et al. [41] noted that digital media lack the ability to be manipulated in 3D, as well as provide multimodal tactile feedback, an important quality of paper documents [37]. They observed a predominance of one-handed interactions in digital tasks, while bimanual interactions were more predominant in the physical tasks. In his Kinematic Chain theory, Guiard observed that in bimanual interaction, the hands coordinate behavior in an asymmetrical manner [12]. We viewed Guiard’s theory as a guiding design principle when dealing with interactions with multiple displays.

Hybrid Piles

Hybrid piles combine the advantages of physical manipulation with the power of digital information. We identify two types of hybrid piles: (1) piles of paper documents, tracked on a digital tabletop or by a video camera, or (2) piles of portable devices containing electronic documents. Khalilbeigi et al. [22] presented interaction techniques for paper piles on tabletops, a good example of the first type of hybrid piles. Their flexible reorganization scheme laid out fluid transitions between neat piles and full juxtaposition, by displaying the documents linearly in a vertical or horizontal matter, or by fanning them out. While we find additional related work in the first hybrid category [19,20,34,35], the second type of hybrid document is the

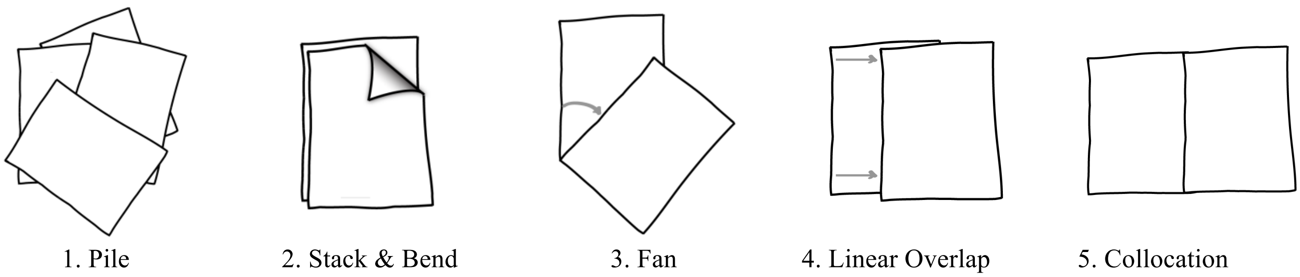

Figure 2. Physical configurations of stacked displays, and selected interaction techniques (gray arrows).

one of interest in this paper. PaperWindows is an early example of these hybrid piles, created through projection on blank sheets of paper [19].

Collocated Displays

While our focuses on the use of stacking, most of the recent work on spatial arrangements of multiple digital documents has focused on the physical collocation of displays. For instance, Chen et al. [7] explored the use of dual display e-book readers, in back-to-back or side-by-side configurations. To turn a page, users flipped the back-to-back display device, while, on the side-by-side display device, they brought the right display towards the left one. Dual displays were found to have the potential to improve the reading experience. Following this work, Hinckley et al. [14] proposed a dual display tablet computer that can be oriented in a variety of postures to support different work contexts.

In addition to placing each document on a different display and arranging them spatially [19], researchers have explored increasing screen real estate by joining multiple collocated computers, or by expanding a single display [21]. This technique is often used to present a single document on two or more screens [6,17] or to extend the desktop to both [15,17,33]. It can also display additional information about a document, such as a broader overview [7]. Finally, to solve the interaction problem faced by dual collocated displays, Hinckley et al. explored connecting two computers through synchronous gestures [17], and created gestures spanning two displays [15].

Piling Flexible Displays

Until now, research on piled displays has mostly focused on simulation, as regular rigid displays were too thick to be stacked with ease. We therefore identified a need to explore electronic documents on piles of thin-film paper-like displays. Emerging flexible display technologies are thin and lightweight, making them ideal for piling and shuffling. While research exists on interaction techniques designed for thin-film displays, implementations did not utilize functional displays. DigitalDesk [43] relied on digital projection on paper, as did PaperWindows [19]. Other work relied on the use of rigid LCD displays on a flexible substrate [36], or paper mockups [28], rather than functional thin-film displays. One research project that did examine interactions with a functional thin-film display was Lahey et al.'s PaperPhone [26]. It investigated the use of bend gestures as a form of input. Results showed users preferred bend gestures that were conceptually simple.

DISPLAYSTACKS OVERVIEW

DisplayStacks (see Figure 1) represents a first step towards a generic e-paper computer interface that implements some paper-like interactions, and in which each computer document is represented by its own paper-like display. It closely resembles DigitalDesk [43] and PaperWindows [19], systems that sought to bring the virtual desktop to the real world, with the distinction that our displays are real, made of Flexible Display Center (FDC) 9.5cm ASU Bloodhound thin-film electro-phoretic displays [35]. While ideally our prototype would have used sheets of 8.5"x11", sheets of this size were not currently technically feasible. Displays are augmented with thin-film bend, location and orientation detection sensors that allow the displays to be aware of their stack location and occlusion by other displays.

BASIC STACKED INTERACTION TECHNIQUES

In DisplayStacks, we explore stacking interaction techniques with multiple displays facing the same direction. Here, we discuss input and interaction techniques through physical configurations created when grouping overlapping displays (Figure 2). We build on traditional configurations and input techniques in the literature, and expand both the basic sets (in this section), and the set of complex interaction techniques (in a later section).

Pile

A pile is a loose grouping of partially overlapping displays (Figure 2.1). A loose set of displays can allow the user to rearrange documents, such as shuffling cards, or pictures. Users can interact with documents in a pile by *moving* the displays among the arrangement, and *inserting* them throughout. This would be the typical way to access and reorder a document in an analog pile.

Stack

A stack is a neat, organized arrangement of displays (Figure 2.2). The user can *insert* displays at different locations within the stack, or *bend* the displays to flick through them.

Fan

A fan configuration is formed by a set of displays shaped in a partial circle pattern (Figure 2.3). It is similar to how players hold a set of cards in their hands. In addition to *insertion*, the user can *rotate*, or *fan*, a display along the fan's pivot, formed typically by the stack's bottom corner. The fan as a display was explored briefly by Lee and Hudson [27], and fanned piles by Khalilbeigi et al. [22].

Linear Overlap

Displays can be arranged in a linear pattern, with partial overlap (Figure 2.4). Steimle et al. [38] refer to this representation as Spread-out. The user can *cover* or *uncover* the displays to interact with them by sliding the displays along the linear configuration, as well as *insert* them as necessary. The uncover interaction technique increases the visible part of the bottom display.

Collocation

Displays can also be side-by-side, collocated [14,17,19], when sharing one edge (Figure 2.5). The user can interact with the displays by *collocating* them.

Transitions and Combinations

It is interesting to note that the user can move through the physical configurations by covering and uncovering displays. Starting in stack mode, the user moves to the horizontal linear overlap by uncovering the top display, and ends up with collocated displays. Khalilbeigi et al. [22] identified these fluid transitions in hybrid paper-and-electronic documents on tabletops. We expanded this to hybrid stacks of electronic documents on flexible displays.

Physical configurations can be combined to create more meaningful interactions with stacked displays. While the illustrations in Figure 2 only make use of two displays, we imagine most interactions with stacked displays will use a minimum of three or four displays [24].

DESIGN RATIONALE

The goal of our design process was to develop tools for organizing digital documents using physical piles and stacks. We identify a number of functionalities typically associated with stacking physical documents that serve as metaphors for the development of complex interaction techniques and applications that ease:

Contextual Overview

Partial stacking can make it easier to get an overview of multiple documents at a time. This is particularly apparent in games of cards, where fanning behaviors allow a user to hold multiple documents, while the content of each display remains identifiable.

Organizing and Sorting

Stacks and piles can contain documents organized according to some parameter, for example time. By stacking incoming documents on a desk, workers can use relative location in the stack to retrieve documents by date. However, traditional piles can be hard to sort. We aim to merge the ability to determine order through relative location in the pile with the ability to sort documents automatically.

Layering Information

In cell based animation, stacks of translucent sheets make it easier to work with layered information. By physically shifting documents between layers, information can be moved relative to that of other documents. We use the layer ordering idea, and borrow from translucent sheets by digitally showing information from displays below.

Nonlinear Browsing

Stacking can ease casual browsing through documents. An extreme example of this is a book, a perfectly neat stack. However, even when documents are not stapled together, being able to insert the finger into a pile and pull up a location in the document allows for physical random access based on location in the pile.

Partial Viewing

The flexibility of paper documents allows for partial revealing of individual documents and partial stacks as a means for browsing the content of the pile without necessarily shifting the order of, or moving individual documents.

Real Estate Increase

In magazines, centerfolds or fold-out spreads allow an increase of the real estate available. Instead of being restricted to images that fit on a single document, the image can be split and displayed on multiple pieces of paper.

Contextual Information

The content of paper documents can support that of other documents. Magazine inserts provide additional information, or half page covers display the headlines, while the full page cover contains the name of the magazine and the photography of a personality.

Design Constraints

We followed a number of design principles related to the actual physical relationship between display and user:

Bending as an Input Metaphor

While flexibility of the display was not a primary concern, we were inspired by some of the bending techniques used in physical document stacks to access information on hidden sheets, and used these as metaphors throughout our design process. Firstly, when browsing a paper stack linearly, users often bend the top corner of a paper sheet to reveal or partially reveal information on the sheet below it. Secondly, when browsing non-linearly through a book, users often bend multiple sheets of paper, for example, to go to a next chapter or index. Thirdly, we were inspired by the use of bends to sort decks of cards, and of dog-earing multiple pages together as a means of binding sorted stacks, and used it as a metaphor for sorting stacked sheets. Finally, we used the metaphor of lifting a transparency sheet from a stack as a metaphor for adjusting transparency in layered documents. Although touch interactions were not technically feasible, our prototype does include stylus interactions.

Two-handed use

We designed our interaction techniques such that they follow patterns of two-handed interactions observed in using paper documents. Not only can bimanual interactions be more efficient, two hands are often *required* when handling multiple documents [12,16]: one to hold the stack, the other to handle individual documents in the stack. Specifically,

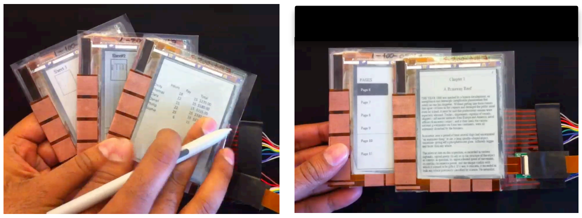

Figure 3. Contextual overview of spreadsheet (left), Contextual menu in a book reader (right).

the non-dominant, generally left hand typically holds the stacks, while the dominant hand interacts with individual displays, e.g. through a stylus, a behavior consistent with Guiard's kinematic chain theory [12].

Thinness and Weight

The main feature of flexible displays that make them suitable for stacking, however, is not their flexibility, but rather, their thinness and low weight. Stacks of thick displays are difficult to hold with one hand, like a stack of cards. We designed our system such that it would potentially be easy to carry the stacks in a pocket. This meant the displays needed to not just be thin, but also as lightweight as possible. Finally, we believe weight plays a role in the dexterity with which individual displays can be moved within and between stacks.

Type of Display

We believed it important to use displays that resemble the reflective properties of paper. While flexible E Ink has the disadvantage of having a relatively slow refresh rate of over 250ms, it resembles physical paper documents more than any other available display technology. Given limited availability, quality and life expectancy of FOLEDs, we chose to work with available Arizona State University (ASU) Flexible Display Center (FDC) 9.5cm Bloodhound flexible electrophoretic displays [35].

Screen Size and Ergonomics

Being unable to explore full size interactions with 8.5"x11" sheets, we focused on designing a display that easily allows holding a stack with one hand, while keeping within the size limitations of current flexible display technologies. We selected a 9.5cm screen size, with all circuitry folded underneath the display to limit the bezel, allowing a stack of at least three displays to be easily held with one hand. The compound displays measure 12.5cm diagonally with bezel.

USING DISPLAYS AS TOOLS

Our basic set of interaction techniques can be combined to form more complex interactions informed by our design rationale. We use displays both for content, but also as contextual tools to manipulate content located on other displays. These include:

Contextual Overviews

To take advantage of the partial screen available in fanned displays, our system allows the user to see a contextual overview of each document in the stack, located on the visible screen portion. Contextual overviews give enough information to identify the content of the document without viewing it in its entirety, exploiting the visual features of the document [20]. In a game of cards, when a hand is fanned, the player sees a "summary" of each card—a colored letter or number, and an image of the suit. Hence, to improve the identification of each display, we digitally augment this action by showing a contextual overview of each document, such as a thumbnail, or a variable from that document, such as the title, date, or version. When working with a spreadsheet, fanned displays give the user an overview of the whole document, with each tab on a different display (Figure 3, left).

Contextual Menus

Fan

One of the problems users face working with single display devices is that valuable screen real estate is often occupied by tool palettes and menus. Contextual menus allow one display to contain contextual tools that can be applied to a second screen. For example, in a painting application, the top display may hold the canvas, while the bottom display holds a palette of paint tools, and the middle display a menu of brush icons. Fanning the displays allows users to increase and decrease the number of tools available. As each display detects what part of the screen real estate is covered, different tools and brushes are displayed in the exposed space, according to frequency of use. A user can dip his or her Wacom stylus onto the second display to select a tool or brush that can then be applied to the top display to draw.

Linear Overlap

Similarly, contextual menus can be popped up across displays by uncovering a display in the stack to a linear overlap position. This serves as a focus + context feature in which the focus may remain on the top document, while a secondary display serves to provide context surrounding that document. For example, when reading a book on the first display, the second display can pop out to the left to show a thumbnail overview of the chapters or pages in the document (Figure 3, right). Pages can be selected from the menu using a stylus, after which the front display shows the selected content. We apply contextual menus proposed by Chen et al. [7] and Khalilbeigi et al. [21] to linear overlap.

Linear Browsing of Piled Information

Piles are stacks that are not neatly organized. In piles, both the location and orientation of displays may be different from display to display (Figure 2.1). When users view information on piled displays, not only may information on a display lower in the stack be partially covered by a higher display, they may also need to reorient displays prior to being able to view the information correctly. Our system allows users to browse information located on lower displays on the top display, in a linear fashion. Bending the top corner of the top display towards the user causes content to cycle through the pile of displays. This interaction technique is inspired by how people go through piles of documents to find the relevant one [34], as well as by Rotated Windows [5]. In a three display pile, upon a bend, information on the middle display is loaded onto the top display, information from the bottom display is loaded onto the middle display, and information on the top display is moved to the bottom display. This allows for a linear browsing action common to piles of unsorted information, like photographs, such that the graphics are always displayed on the top display in the proper orientation. Piles may be held with the non-dominant hand or placed on a surface.

Non-Linear Browsing & Sorting of Stacked Information

One of the problems with stacks is that they can be tedious to sort [23]. Our system handles sorting automatically by transferring information between displays according to some variable, such as date and time of creation, or alphabetically. A bend of the top right corner of the entire stack away from the user sorts the information on the displays according to the first variable, a second bend according to the second variable, and so forth.

For instance, a set of displays that represent electronic business cards picked up at a conference may be placed in a pile, then arranged in a stack. Bending the set of displays away from the user causes the business cards to be sorted by last name. Bending it again causes them to be sorted by date and time at which the card was received. Bending it again sorts the cards according to nationality, while a final bend renders them unsorted again. Business cards can be browsed linearly, i.e., in order of display, by bending the top right corner of the top display towards the user.

A second problem of stacks is that displays obscure each other. In the physical world, documents need to be pulled out in order to be viewed. Our system allows bends of multiple displays towards the user to actuate non-linear browsing. Linear browsing is actuated by bending the top right corner of the top display towards the user. For example, if each display contains a chapter of a book, bending the first display will page through the first chapter. With non-linear browsing, the user bends multiple screens, e.g., the first, second and third display simultaneously. This places the third chapter onto the top screen. This bend swaps the two chapters: the first chapter is now displayed on the third screen. Single bends of the top screen now allow paging through the third chapter. A second bend of the first, second and third display causes the chapters to return to their original order.

Merging and Splitting Documents

Documents that span multiple displays in the stack can be merged into a single document on the top display by bending the left side away from the user. The bottom displays become blank after the merge operation. Individual pages in a document displayed on the top display can be split into separate documents on the other (blank) displays in the stack by bending the left side of the stack towards the user.

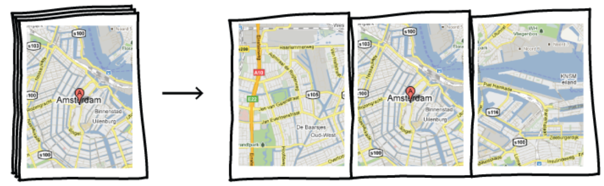

Figure 4. By collocating displays, we increase the real estate to display a wider map.

Layering and Transparency

With physical stacks, the use of transparencies (e.g., overhead projector slides) allows for see-through effects that have proved useful in cell-based animation, photography, architecture and medicine. Stacks of thin-film displays offer two main advantages: to control 1) the digital transparency level of images, and 2) the position and orientation of layers to create aligned graphics. By using physical stacks, every layer is physically represented, and can be pulled out of the stack in order to be examined in isolation, or to be reordered.

While in the stack, the level of transparency of individual layers (i.e., displays) can be adjusted by bending its bottom right corner of the corresponding display in the stack, with an upwards bend increasing the level of transparency. The transparency level of groups of displays can be altered by bending the bottom right corner of the set of displays. Alternatively, users can use a fanned display menu that provides contextual menus that control transparency. To create a specific composition, the user can apply linear and non-linear browsing as well as sorting techniques to the layers through bends of the top right corner of the displays. Once an appropriate layering of information has been achieved, layers can be flattened (merged) by bending the left side away from the user, leaving the user with a single display that shows the composite outcome.

Increasing Screen Real Estate through Collocation

One way of solving the problem of limited screen real estate is to collocate two separate displays such that they form one single, larger screen [15,17]. Linear overlaps can be used in this process to continuously widen a display to the exact size needed. Once the displays are collocated, applications adjust automatically to the larger canvas. While browsing a map on a stack of three displays, the area shown on the map can be enlarged to three times its size by uncovering one display to the left, and another to the right of the central display (Figure 4). Another example is in an increased sketching canvas space. Note that our system currently has the disadvantage of introducing visible bezels that obscure part of the display.

Physical Sorting and Insertion

The stack of digital displays retains all the benefits of being able to physically sort the displays by popping them in or out of the stack. Physical tying of information to a particular display, however, is overridden with any of the interaction techniques that move information between displays.

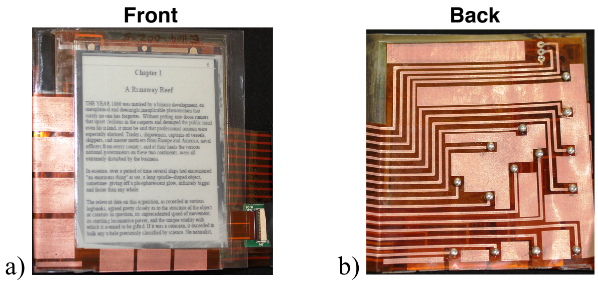

Figure 5. Detecting relative location of displays using a conductive dot pattern between the front a) and back b).

This means there is no strong tie between information and any one particular display.

DISPLAYSTACKS IMPLEMENTATION

To implement the interaction techniques described, we augmented a number of flexible E Ink displays with thin-film stack sensing technology. The DisplayStacks prototype tracks the relative position of the flexible display screens in a physical stack as well as the dynamic movement/repositioning of the displays within the stack. Based on the interactions by the user, the displays are updated to reflect changing states of location and orientation in the stack. Within a pile, our display regions change dynamically to present context-specific content. To our knowledge, no prior work exists that has demonstrated these techniques with flexible thin-film displays.

We created the system using 3 displays, but DisplayStacks is scalable to as many flexible displays as are needed for a particular application. We believe a plurality, i.e. 3 displays, would suffice to explain most techniques, such as bending the corner of a single display for non-linear browsing, or of the entire stack for sorting, which all scale up. Kim et al. [24] noted that stacks can be as small as 3 or 4 displays.

Prior Work on Tracking Piles & Collocated Documents

Various systems have been developed to track interactions with physical piles of objects. Previous works on tracking piles of paper and books have used radio-frequency identification tags [4], computer vision [11,24], motion capture [19] and conductive inks [32]. While these systems explore interesting approaches to tracking objects within a pile, they either require specialized desk setups, including cameras [11,19,24], or they limit the tracking to a fixed order of documents [4,32]. ConnecTable [40] used passive tags based on radio frequency transponder technology. Siftables [33] detect their own motion with accelerometers, and their proximity through infrared transceivers. Other methods include bumping displays detected through accelerometers [17], and ultrasound [25]. Hoffman et al. suggested the electrical contact sensing that was used in our system [18].

Prototype

To implement our interaction techniques, our DisplayStacks prototype uses Arizona State University Flexible Display Center (FDC) 9.5cm Bloodhound thin-film electro-phoretic displays [35]. Figure 5 shows a flexible display augmented with thin-film bend and location detection sensors. The total thickness of a display is 3 *mm*. DisplayStacks is currently wired to rigid electronics driving the flexible displays. The displays are given mobility through long ribbon cables.

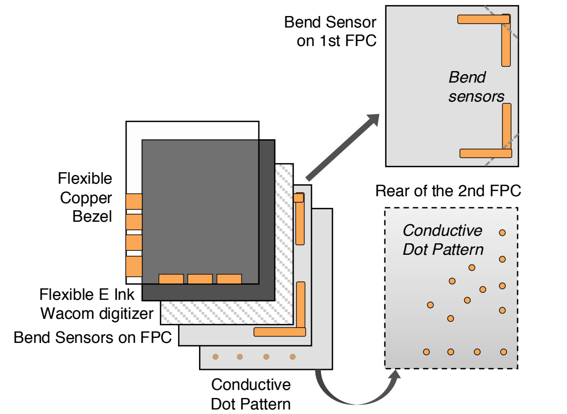

Figure 6. Each display is augmented with 4 layers: a conductive bezel on a flexible printed circuit (FPC), 4 bend sensors on an FPC, a Wacom digitizer, and a conductive dot pattern on an FPC.

Each flexible display is augmented with four layers of sensors (see Figure 6). The first layer is located along the bottom and left bezel of the display, where the non-dominant hand typically holds the stack, ensuring better connectivity between layers. It consists of seven flexible conductive zones that detect the order in which displays are stacked, laid out on a flexible printed circuit (FPC). The second layer, beneath the display, is a flexible Wacom [42] digitizer. The third layer has 4 Flexpoint 2” bi-directional bend sensors [10], on an FPC. The fourth layer consists of an asymmetrical conductive dot pattern, with 14 conductive dots, on an FPC. As shown in Figure 5b, the dot pattern faces outwards on the back of the display. The dot pattern interacts with the top bezel (Fig. 5a) to determine location and orientation of the displays within a stack. Each dot is 1 *mm* thick, protruding from the display to ensure a properly conductive connection with the bezel of the display below. The two FPCs were built by printing our circuit design onto a sheet of Dupont Pyralux copper coated polyamide using a solid wax printer. To produce the circuitry, the Pyralux was etched using hydrochloric acid, after which the wax was dissolved to expose the circuitry.

Processing

The flexible display and the Wacom digitizer are driven by an E Ink [9] Broadsheet AM300 Kit featuring a Gumstix [13] processor. An Arduino Mega 2560 microcontroller [3] obtains data from the other three sensor layers. Displays and sensors interact through a Max 5 [31] patch that processes sensor data from the Arduino, runs the logic to determine the current state of each flexible display, and updates the flexible displays via the AM300s.

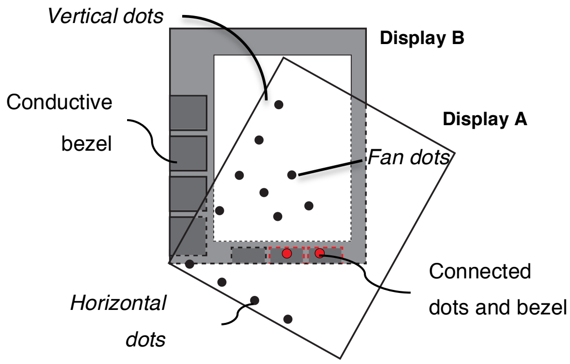

Figure 7. A subset of the conductive dots of the top display forms a circuit with a subset of the conductive bezel of the bottom display (illustrated in red). This forms a fanned occlusion. Dots along the horizontal and vertical axes are connected to the respective horizontal and vertical conductive bezels when moving displays horizontally or vertically.

Recognizing Interaction Techniques

Interaction techniques are implemented by interpreting data from the four sensor layers. When two displays are stacked, the conductive dot patterns on the back of the first display connect with the corresponding conductive bezel of the display below. The relative location and orientation of the displays is based on which dots connect to the bezel. The bend sensor layer enables interaction with the stack using a rich set of bend gestures. A Wacom digitizer allows pen-based interactions with individual displays in the stack.

Sensing Location and Orientation

When one display (display A) is stacked on top of another display (display B), a subset of the conductive dot pattern beneath display A makes contact with a subset of the conductive zones on the front of display B to form a circuit (see Figure 7). When a conductive dot connects with a conductive zone, this is registered by the Arduino and sent to the Max program. It keeps track of the activated dot-bezel circuits, thus reconstructing the relative position and orientation of display A with regards to display B. When a third display is added below display B, its relative position and orientation can be determined as well. Stack order is determined by sensing a unique electrical signal sent by each display over its bezels.

Detecting Occlusions

When a display is moved in a stack, the system uses its location and position data to compute the size of the area of display B that is not occluded by display A. Depending on the state of the display stack, this information is then used to update the graphics on the display. We selected a specific dot pattern (Figure 7) to support stacking, discrete steps of linear overlap along the horizontal and vertical axes and fanning out. Due to the limited resolution of the dot patterns, each display currently detects eight zones of occlusion: 3 partial vertical occlusions, 2 partial horizontal occlusion, and 3 fanned occlusions. The patterns are divided in three zones to recognize three types of occlusion. Dots along the vertical axis, on the side of the display, recognize vertical occlusions when they come in contact with the left, vertical bezel. Dots along the horizontal axis, on the bottom of the display, recognize horizontal occlusions when they come in contact with the bottom, horizontal bezel. Fanned occlusions are recognized when the dots in the middle of the screen make contact with the bottom bezel (as illustrated in Figure 7). Stacking is detected when vertical dots come in contact with the left bezel simultaneously with horizontal dots coming in contact with the bottom bezel.

The system is robust enough to track 8 different configurations, by maintaining a state machine for each display. The displays start in the *unstacked* state. In this state, displays can either not be overlapping, or loosely piled. Once aligned, they become a *stack*. In this state, the user can thumb through the displays, or insert displays in the stack. The user can also fan out the displays to move to a *fan* state, uncover the displays horizontally to move to the *linear overlap (horizontal)* state, or uncover vertically to move to the *linear overlap (vertical)* state. In each *linear overlap* state, the user can cover and uncover the displays. Once completely uncovered, the state becomes *collocation*. Specifically, collocation is currently detected by the change from *linear overlap* to *no contact* state. In every state, the user can align the displays to return to the *stack*, or loosen the displays to return to the *unstacked* state. While we use redundant dots (e.g. 2 dots can identify the fanned state in Figure 7) and state-tracking to improve reliability and reduce false positives, our simplified sensing is limited in that it does allow the possibility of false positives between *collocation* and *unstacked*, something investigated in our initial user study.

Detecting Bends

We used 4 bi-directional bend sensors, located at the top and bottom right corners, as shown in Figure 6. Users can perform 6 bends with the two pairs of bend sensors embedded in the prototype: top right corner up/down; bottom right corner up/down, and right center up/down. Perpendicular pairs of sensors are used to obtain an optimal measure of the flexion of each corner using redundancy.

Initial User Experiences

We tested the accuracy of the sensing technology by asking 6 users to achieve each of the 8 configurations from the stacked state (repeated 5 times). Preliminary results indicated that the system correctly recognized states with about 75% accuracy. Users preferred the vertical overlap overall (82% accuracy), with the horizontal overlap coming close second (77% accuracy). Fanning was not a preferred technique, as it involved more effort than the other techniques (66% accuracy). However, two users learned to pivot the displays by holding the stack with their thumbs, allowing them to execute fanning with the same effectiveness as the other techniques.

DISCUSSION

With DisplayStacks, we proposed a set of interaction techniques for electronic documents on thin-film flexible E Ink displays that enable the physical manipulation of electronic documents in ways similar to paper documents. We believe there are several key benefits to this:

Physicality of Stackable Windowed Content

One of the chief benefits of flexible displays is that they can be integrated into a very thin form factor. This allows user manipulations that are currently not available with LCD-based computing devices. One clear benefit of the use of many thin-film displays in a single workspace is that windowing no longer serves as the primary means of managing workflow between multiple documents views. Instead, each display correlates with a window. Because the displays are thin, windows can be stacked similarly to electronic windows. We argue that this stacking behavior brings organizational benefits that are difficult to implement using multiple LCD-based displays, if only because their thickness makes it difficult to handle stacks and fans of multiple displays.

Tactile Representation of Windowed Documents

DisplayStacks' interaction techniques provide tactile feedback in ways that traditional displays do not. Holding stacked displays gives the user tactile information about the total number of displays or windowed documents available. Bending as an input technique allows for a physical correlate that represents the action directly in the muscle receptors of the user [2]. This can help users achieve light-weight interactions without requiring visual attention. Bending multiple screens allows for actions such as sorting to be applied easily across arbitrary groups of (obscured) display devices, as well as be tangibly represented.

Dynamic Display Regions

Within a pile, our display regions change dynamically to present context-specific content. To our knowledge, no prior work exists that has demonstrated these techniques with flexible thin-film displays. Previous work has discussed collocation with rigid, thick LCD displays [15,17], as well as introduced bookmarks when working with dual E Ink screens [7]. We believe that the thickness of the stacked displays to be stacked represents a critical design parameter. With DisplayStacks, we expanded the work and presented a number of new configurations (e.g. fan), input techniques (e.g. rotating displays, bending single display, bending piles), and interaction techniques (e.g. contextual overviews, menus, layering, merging documents). These new interaction techniques all make use of either dynamic display regions, or take advantage of the thinness and flexibility of the displays. Note that dynamic updates of display contents may, however, lead to confusion amongst users during the re-orientation or re-positioning of a display.

Permanence of Information

As we compare the interaction techniques on thin-film displays to those of paper, it is critical to discuss the permanence of information. The information displayed on paper is highly permanent, while the content of a specific display is not. The fixity of content with respect to its medium can help in ad hoc sorting tasks [29], aiding recall for where information is located on a desk, stack or page [23]. Note that an important benefit of E Ink displays is that they hold their contents even when they are not powered. This means that displays, and thus windowed documents represented by these displays, maintain the state of their content even when not in use, which represents a form of fixity of content.

Strengths and Limitations of the Current Prototype

Two main strengths set apart the conductive dot pattern technique from other systems tracking piles: the absence of a specialized desk setup, and the ability to track documents in any order and position. While many prior systems are vision-based [11,19,24], or tabletop-based [22,38,39], DisplayStacks provides tracking hardware both independent from the desk and self-contained in each document. In addition, it can track the movements of one or multiple displays at once, in any order, an improvement on previous work [4,24,32]. While it does require specialized circuitry on each display, the addition of those components is reasonable at this time as we have not reached ubiquity of flexible displays.

However, the conductive dot pattern tracking needs to be more accurate, limits gestures to be discrete, and limits the number of orientations and positions available. A prefabricated circuit with a denser concentration of conductive dots and zones could provide finer grained information that computes the position and orientation of displays in a stack more precisely. This could allow for continuous rather than discrete occlusion detection. Finally, although three displays are enough to demonstrate each interaction technique, the use of a limited number of displays restricts combinations of interaction techniques, as well as may not reflect real world piles. Another core limitation of this work resides in the physical restrictions of the prototype, mainly based on restrictions in the display technology. The design of the flexible display requires the presence of a bezel to power the pixels, leaving a gap between displays during overlap or collocation. This impacts the illusion of increased real estate. In addition, current flexible display prototypes cannot be bent everywhere, due to the presence of rigid electronics. This reduces the pallet of bend gestures available for consideration. Another critical restriction is that our displays are currently tethered for technical reasons. We hope to realize future versions of the system in which electronics and batteries are contained in a flexible sheet of material.

FUTURE WORK AND APPLICATIONS

Future applications of an untethered version of DisplayStacks exist, e.g., in mobile scenarios. By carrying a stack of displays, users can increase their real estate on the fly, e.g., when studying maps [21]. Each display may run its own mobile app, easing multitasking through the use of multiple lightweight displays. E.g., a user might use one display to make notes while holding a videoconference on a fanned display. Interactions between screens can facilitate a lightweight means of moving data between apps or windows, for example, to send notes as an attachment to an email may only require tapping the note display onto the display with the reply email. In the future, we therefore need to explore other techniques for creating, transferring and deleting content on e-paper computing interfaces. Kim et al [24] proposed a video-based document tracking system that remembers the location of each document in piles. We would like to extend this work to study the effect and interaction techniques related to grouping. Finally, we look forward to performing a formal user evaluation of our system.

CONCLUSION

We presented DisplayStacks, a system that allows users to interact with stacks of multiple thin-film flexible displays containing digital documents. DisplayStacks introduces tools for organizing *digital* documents using piles and stacks of *physical* display windows. DisplayStacks is a functional prototype composed of multiple thin-film E Ink electronic paper displays augmented with conductive dot pattern sensors that detect relative position and orientation of displays within the stack. With this system, we provide users manipulating hybrid stacks of windows with the ability to combine the benefits of physical manipulation of paper documents with the malleability of electronic content.

REFERENCES

1. Agarawala, A. and Balakrishnan, R. Keepin' It Real: Pushing the Desktop Metaphor with Physics , Piles and the Pen. *Proc. CHI*, (2006), 1283-1292.
2. Akamatsu, M., MacKenzie, I.S., and Hasbrouc, T. A comparison of tactile, auditory, and visual feedback in a pointing task using a mouse-type device. *Ergonomics 38*, (1995), 816-827.
3. Arduino. . http://www.arduino.cc.
4. Back, M.J. and Cohen, J. Page detection using embedded tags. *Proc. UIST 2*, (2000), 159-160.
5. Beaudouin-Lafon, M. Novel interaction techniques for overlapping windows. *Proc. UIST 3*, 2 (2001), 153.
6. Borovoy, R. and Knep, B. Junkyard Jumbotron. 2011. http://jumbotron.media.mit.edu.
7. Chen, N., Guimbretiere, F., Dixon, M., Lewis, C., and Agrawala, M. Navigation techniques for dual-display e-book readers. *Proc. CHI*, (2008), 1779–1788.
8. Davidson, P.L. and Han, J.Y. Extending 2D object arrangement with pressure-sensitive layering cues. *Proc. UIST*, (2008), 87.
9. E Ink Inc. . http://www.eink.com.
10. FlexPoint Inc. http://www.flexpoint.com
11. Fujii, K., Shimamura, J., Arakawa, K., and Arikawa, T. Tangible search for stacked objects. *Proc. CHI extended abstracts*, (2003), 848.
12. Guiard, Y. Asymmetric Division of Labor in Human Skilled Bimanual Action: The Kinematic Chain as a Model. *J Mot Behav 19*, 4 (1987), 486-517.
13. Gumstix Inc. http://www.gumstix.com
14. Hinckley, K., Dixon, M., Sarin, R., Guimbretiere, F., and Balakrishnan, R. Codex: a dual screen tablet computer. *Proc. CHI*, (2009), 1933–1942.
15. Hinckley, K., Ramos, G., Guimbretiere, F., Baudisch, P., and Smith, M. Stitching: pen gestures that span multiple displays. *Proc. AVI*, (2004), 23–31.
16. Hinckley, K., Yatani, K., Pahud, M., et al. Pen + Touch = New Tools. *Proc. UIST*, (2010), 27-36.
17. Hinckley, K. Synchronous gestures for multiple persons and computers. *Proc. UIST*, (2003), 149-158.
18. Hoffmann, F. and Scott, J. Location of Mobile Devices Using Networked Surfaces. *Proc. Ubicomp*, (2002), 281-298.
19. Holman, D., Vertegaal, R., Altosaar, M., Troje, N., and Johns, D. PaperWindows : Interaction Techniques for Digital Paper. *Proc. CHI*, (2005), 591-599.
20. Jakobsen, M.R. and Hornbæk, K. Piles, tabs and overlaps in navigation among documents. *Proc. NordiCHI*, (2010), 246–255.
21. Khalilbeigi, M., Lissermann, R., Mühlhäuser, M., and Steimle, J. Xpaaand : Interaction Techniques for Rollable Displays. *Proc. CHI*, (2011), 2729-2732.
22. Khalilbeigi, M., Steimle, J., and Mühlhäuser, M. Interaction techniques for hybrid piles of documents on interactive tabletops. *Proc. CHI extended abstracts*, (2010), 3943.
23. Kidd, A. The marks are on the knowledge worker. *Proc. CHI*, (1994), 186–191.
24. Kim, J., Seitz, S.M., and Agrawala, M. Video-based document tracking: unifying your physical and electronic desktops. *Proc. UIST*, (2004), 24–27.
25. Kortuem, G., Kray, C., and Gellersen, H. Sensing and visualizing spatial relations of mobile devices. *Proc. UIST*, (2005), 93.
26. Lahey, B., Girouard, A., Burleson, W., and Vertegaal, R. PaperPhone : Understanding the Use of Bend Gestures in Mobile Devices with Flexible Electronic Paper Displays. *Proc. CHI*, (2011).
27. Lee, J.C., Hudson, S.E., and Tse, E. Foldable interactive displays. *Proc. UIST*, (2008), 287.
28. Lee, S.-S., Kim, S., Jin, B., et al. How users manipulate deformable displays as input devices. *Proc. CHI*, (2010), 1647.
29. Malone, T.W. How do people organize their desks?: Implications for the design of office information systems. *ACM Transactions on Office Information Systems 1*, 1 (1983), 99-112.
30. Mander, R., Salomon, G., and Wong, Y.Y. A "pile" metaphor for supporting casual organization of information. *Proc. CHI*, (1992), 627-634.
31. Max 5. http://www.cycling74.com
32. May, K.W. Conductive Ink Based Page Detection for Linking Digital and Physical Pages. *Proc. CHI*, (2001), 275-276.
33. Merrill, D., Kalanithi, J., and Maes, P. Siftables: towards sensor network user interfaces. *Proc. TEI*, (2007), 75–78.
34. O'Hara, K. and Sellen, A.J. A Comparison of Reading Paper and On-Line Documents. *Proc. CHI*, (1997), 335-342.
35. Raupp, G., O'Rourke, S.M., Moyer, C., et al. Low-temperature amorphous-silicon backplane technology development for flexible displays in a manufacturing pilot-line environment. *Journal of SID 15*, 7 (2007), 445.
36. Schwesig, C., Poupyrev, I., and Mori, E. Gummi: a bendable computer. *Proc. CHI*, (2004), 263 - 270.
37. Sellen, A.J. and Harper, R.H.R. *The Myth of the Paperless Office*. MIT Press, 2003.
38. Steimle, J., Khalilbeigi, M., Mühlhäuser, M., and Hollan, J.D. Physical and Digital Media Usage Patterns on Interactive Tabletop Surfaces. *Proc. ITS*, (2010), 167-176.
39. Steimle, J., Khalilbeigi, M., and Mühlhäuser, M. Hybrid groups of printed and digital documents on tabletops. *Proc. CHI extended abstracts*, (2010), 3271.
40. Tandler, P., Prante, T., Muller-Tomfelde, C., Streitz, N., and Steinmetz, R. Connectables: dynamic coupling of displays for the flexible creation of shared workspaces. *Proc. UIST*, (2001), 11–20.
41. Terrenghi, L., Kirk, D., Sellen, A.J., and Izadi, S. Affordances for manipulation of physical versus digital media on interactive surfaces. *Proc. CHI*, (2007), 1157–1166.
42. Wacom Inc. http://www.wacom.com
43. Wellner, P. Interacting with paper on the DigitalDesk. *Communications of the ACM 36*, 7 (1993), 87–96.
44. Whittaker, S. and Hirschberg, J. The character, value, and management of personal paper archives. *TOCHI 8*, 2 (2001), 150-170.