

PaperTab: Windowing Techniques for an Electronic Paper Computer

Aneesh Tarun¹, Peng Wang¹, Audrey Girouard¹, Paul Strohmeier¹, Derek Reilly² and Roel Vertegaal¹

¹Human Media Lab, Queen's University,
Kingston, Ontario, Canada

²Faculty of Computer Science, Dalhousie University,
Halifax, Nova Scotia, Canada

ABSTRACT

We present PaperTab, a paper computer with 9 Letter-sized functional touch sensitive flexible electrophoretic displays. PaperTab merges the benefits of working with digital information with the tangibility of paper documents. In PaperTab, each document window is represented as a physical, functional, flexible e-paper screen called a tab. Each tab can show documents and other digital information at varying resolutions. The location of tabs is tracked on the desk using an electromagnetic tracker. This allows for context-aware operations between tabs and the desk. The desk is divided into 3 proximity-based zones: the hot zone, used for full screen browsing and editing of documents, the warm zone for displaying thumbnail overviews, and the cold zone for filing of documents. Touch and bend sensors in each tab allow users to navigate content. Tabs can also be pointed at one another for focus+context view operations: e.g., documents are opened by pointing a tab at a file icon on a second tab. Bend operations on a tab allow users to navigate content: flicking the top right corner pages forward or back, while bending the tab zooms in or out of the document. We report on a user experience study with 12 participants.

Author Keywords

Flexible Electrophoretic Displays, Physical Windows, Paper Computers, Organic User Interfaces.

ACM Classification Keywords

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

General Terms

Design, Human Factors.

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




Figure 1. PaperTab with 9 physical windows and (virtual) hot, warm and cold zones.

INTRODUCTION

A long-standing debate in user interface research is the tradeoff of benefits between physical and digital user interface objects. In particular, the vision of a physical desktop computing system based on the way office workers use paper documents has been an enduring research goal [5]. One of the reasons for the longevity of paper, according to Sellen and Harper [30], is that it provides tactile-kinesthetic feedback when organizing and navigating information that is not available in traditional digital windowing environments. Paper, as a physical medium, is also thin, lightweight and portable. It provides 3D spatial organization of information, while enabling concurrent access to multiple streams of information [30]. On the other hand, Graphical User Interfaces (GUIs) provide superior opportunities for on-the-fly electronic manipulation and updating of information over paper. While accepting that malleability is a basic requirement for any user interface, in this paper, we address three major limitations of the GUI, as compared to paper documents: (1) users are severely restricted in the way they concurrently manipulate and organize multiple windows, particularly in cases where windows obscure each other; (2) spatial manipulation of windows is defined and limited by screen size and (3) users cannot apply spatial memorization skills for GUI-based information retrieval as effectively as they can in real, physical, environments [16]. One solution is the design of a system that combines tangible, paper-like interactions with digital information (not limited to documents), specifically by embodying windows onto an electronic paper-like medium.

The idea of providing paper-based interaction with digital information is far from new: it was the core idea behind Memex [5] which inspired the GUI. Wellner's DigitalDesk [34] was one of the first systems to implement such seamless interaction between physical and digital mediums through digital projection on paper. Since then, many tabletop research projects have explored the coexistence of paper and digital information [15,17]. While these systems provide various affordances of paper, they are often limited in their interactions, and do not take advantage of the spatiality of content beyond the 2D surface of a desk.

The emerging technology of thin-film flexible displays [18] presents an opportunity to merge the physical world of paper with the digital world of information via Organic User Interfaces [32]: non-flat, flexible, tactile, high-resolution display interfaces. Flexible displays are sufficiently thin to approximate paper-like interactions, and sufficiently light to allow for efficient spatial interactions between displays. While research on flexible display interfaces has pursued the embodiment of digital data on paper [16,18], this research has mostly focused on the usability of *single* and *small* display interactions. In this paper, we extend this work to multi-display large format interactions. We present PaperTab, an electronic paper computer that allows physical manipulation of digital information using multiple flexible electrophoretic displays embodied on a physical desk (see Figure 1). Documents in PaperTab combine the *malleability* of electronic windows with the tactile-kinesthetic and 3D spatial *manipulability* of paper documents. In PaperTab, each graphical window is represented by a fully functional, paper-sized 10.7" diagonal Plastic Logic thin-film high resolution flexible electrophoretic display [24].

Contributions

This paper contributes one of the first paper computers with many real, functional, and large touch-sensitive thin-film flexible electrophoretic displays. We also contribute a functional physical windowing system with interaction techniques for zone and focus+context interactions. Zone interactions are based on the proximity of each display to the user. PaperTab determines the 6 DOF location and orientation of each display on the desk through an electro-magnetic sensor mounted on the back of the displays. Locations of displays are categorized into hot zones, used for active editing, warm zones for temporary storage, or cold zones for long-term storage [30]. Each display has a transparent flexible capacitive touch input layer. This is used for focus+context interactions [3], which allow users to point with one flexible display at coordinates within a second flexible display, e.g., to pop up a magic lens [4]. We contribute a user experience study, examining several representative tasks in which data was moved within and between multiple flexible displays.

BACKGROUND

Typical window managers and file systems provide capacities similar to the physical desktops they originally mimicked, including the ability to move work in and out of focus

based on its position and relative size on the desktop, to juxtapose and arrange work documents using a 2-D spatial layout and, to some extent, to recall documents based on their location in the 2-D workspace. With PaperTab, we were interested in the extent to which this evolved 2-D environment can be applied back onto the physical desktop, using lightweight flexible displays that are spatially registered on the workspace. This "restoration" of the desktop is not trivial: it differs from the purely digital desktop by employing direct, tangible and bimanual manipulation of physical artifacts, and differs from the purely physical desktop since these artifacts are dynamic and interactive. In this section we review the significant base of related work that has inspired and motivated our design.

Paper Computers

DigitalDesk [34] was one of the first physical paper computers. It seamlessly merged interactions between physical paper and digital documents on a physical desk. Users were able to select data from paper documents and copy it into digital documents. In PaperWindows, Holman et al. [16] created a windowing environment that simulated fully wireless, full-color digital paper. PaperWindows demonstrated use of gestural inputs such as hold, collate, flip and bend. Similarly, Lee et al. [19] used image projection on foldable materials to simulate flexible displays with variable form factors and dimensions.

None of these systems deployed functioning thin-film electronic paper displays. Research in thin-film display interactions started with paper mockups, bendable substrates on rigid devices and projected flexible displays [16,29]. With the recent availability of working flexible displays, projects like PaperPhone [18] explored new interaction techniques such as bending as an interaction paradigm. DisplayStacks [11] is one of the few papers to explore interaction techniques for stacking multiple functional E Ink displays. However, it featured smartphone sized displays, and presented a set of interactions between screens based on stacking metaphors.

Hybrid Documents

Over the past decades researchers have also explored the integration of digital content with a range of physical media, from paper on desks [16,34] to notebooks [21] and maps [25]. A number of relationships between digital and physical media have been studied, including designing physical media explicitly to support querying of digital content [25], or synchronizing digital and physical manifestations of the same document [13]. However, there remain few universally accepted design principles for such integrations [20].

Khalilbeigi et al. [17] worked on a tabletop system that tracked various configurations of physical documents, as did Hinckley et al. [15]. Interaction with hybrid physical-digital documents can occur by pointing at [9], writing on [13] or manipulating [16] physical media, by moving a digital display in relation to physical documents (e.g. in magic lens interactions [4]), or by interacting with the GUI superimposed




Figure 2. PaperTab tab in the warm zone showing thumbnails of a photo album document.

on physical media [16] or on a separate digital device [9]. In several cases the media itself serves as both display and input: e.g., Reilly et al. [25] find that users tend to choose a corner of the media to facilitate precise pointing and to reduce occlusion.

Physical Interaction on a Plane

Interaction paradigms for 3-D virtual environments have tried to exploit egocentric perspective and proprioception [23] for target acquisition. Prior work has illustrated greater effectiveness of a 2-D or 2.5-D workspace (either physical or virtual) over 3-D workspaces [2], in terms of both target acquisition performance [6] and in recall of item locations [7]. Cockburn and McKenzie suggest the superiority of a physical 2-D interface to a virtual one in terms of task efficiency and subjective performance measures [6]. When comparing physical and virtual interfaces, however, Terrenghi et al. [31] observed faster puzzle solving using physical objects, and equal performance for sorting. In addition, they exhibited significantly more bimanual interaction in the physical interface conditions. The authors also observe fundamentally different methods of interaction between physical and virtual tabletop activities, related in part to the ability to hold and place multiple physical artifacts. Their participants also preferred the physical implementation.

Bridging Physical and Virtual Worlds

Token+constraint systems [33] classify physical objects as tokens manipulated under a set of space-based constraints to query data and execute procedures. The basic token in PaperTab is the physical document window, constrained dynamically in terms of its placement to other physical document windows, and absolutely in terms of proximity to the user. This approach—defining constraints dynamically in terms of tokens that are themselves dynamic—has been explored in other systems such as Siftables [22], and Datatiles [27]. PaperTab extends this to include paper-inspired interactions with context-aware physical windows, in which each window is represented by one display.

PAPERTAB OVERVIEW

PaperTab is a paper computer made of many flexible thin-film electrophoretic displays that can be held in the hands or placed on a physical desk. The displays serve as electronic paper windows into computer documents. In PaperTab, each display corresponds in functionality to one window on a GUI. We call these physical instantiations of one window per display a *tab* (see Figure 2). Users can interact with the system by moving tabs around their desk, and can navigate and edit information on a tab using touch input and simple bend operations [18]. Users can also use a tab as a magic lens that points at another tab, allowing them to pop up information on the top tab detailing items displayed on the underlying tab.

DESIGN RATIONALE

In designing PaperTab, we focused on developing context-aware window management techniques that would allow for interactions with many documents represented on many physical displays, through relative movement of those displays. We used the following design criteria:

1 Window = 1 Display

The basic token in PaperTab is the *tab*: each window *is* a physical display. The main reason for this was that it allows for virtually unlimited screen real estate spread over large work areas. PaperTab is document centered, and requires a physical instantiation of each document on a tab. While tabs *can* have multiple documents, and documents on tabs *can* have multiple pages, at no point is a document not represented on a tab. This allows users to move digital documents by moving tabs physically and tangibly through spatial arrangement on their desk.

Location Awareness

Each tab is aware of its own location on the desk, as well as relative to and within other tabs. This allows smart context-aware operations between tabs that depend on how close they are to the user or to other tabs.

Spatial Proximity = Resolution

Because screen real estate is limited in desktop GUIs, windows are typically stacked on the desktop in a virtual *z* dimension, obscuring one another. Since in PaperTab real estate is not limited to a single display, or even a single desktop, we designed tabs to vary the resolution of a document's information based on physical proximity to the user [30]. Tabs that are furthest away from the user represent file thumbnails that are not in use. Documents in the middle zones represent overviews.

Spatial Proximity = Focus

The use of proximity to activate document views also required a redefinition of the concept of top or focus window. Again we used the metaphor of proximity: the closer a display is to the user, the more focus we expect it to receive and the more receptive to input it becomes. According to Sellen and Harper [30], users often use proximity when dealing with paper documents. *Hot* paper documents are in immediate use, serve multiple concurrent purposes and are in close

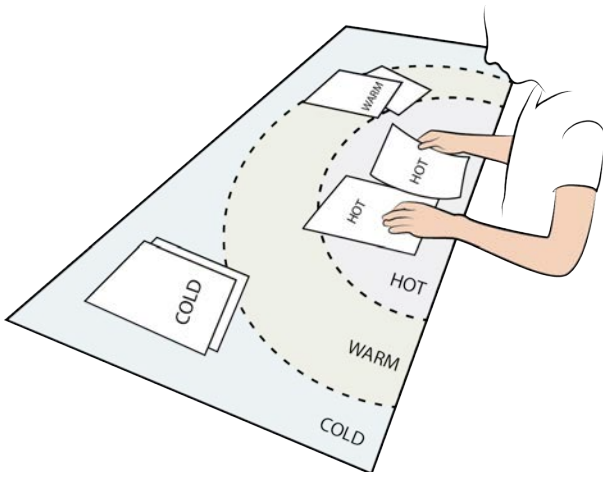


Figure 3. Hot, warm and cold proximity zones in PaperTab.

proximity to the user. *Warm* paper documents have served or are about to serve an immediate need, but should not be in the way of handling hot documents. *Cold* paper documents represent archived documents that are not in immediate use, and are typically filed away from the user. We translated these notions of hot, warm, and cold zones into a proxemics design [12] in which tabs are brought in an out of task focus by physically moving them towards or away from the user. Sellen and Harper [30] also suggest that to allow for easy multitasking, every document that is within close proximity of the user should be active for input, with the most active windows being the ones held by the user.

Moveability = Mass and Volume

The effectiveness of having displays as windows is greatly reduced by the weight of the display [8]. Moving a stack of 10 tablets around one's desk does not rival the efficiency of moving paper documents. To move tabs around with the greatest efficiency, we believe it is important that they are made of the most lightweight displays available.

Display = Pointing Device

We were greatly inspired by work on *focus+context* displays [3], magic lenses [4] and spatially aware computing [9] in the design of our system. As with lenses in GUIs, users should be able to use one tab to navigate content displayed on another tab.

INTERACTION TECHNIQUES

In this paper, we focused the design of interaction techniques for PaperTab on the problem of windowing in a system with many displays, where each display is the functional equivalent of a single GUI window. One challenge was that unlike GUIs, no user interface exists in the space between windows. This means the user interface lives among a collection of displays, rather than just within a single display. As such, we considered interactions with tabs as pertaining to one of two types: zone interactions and focus+context interactions.

Zone Interactions

In traditional GUIs, windows are organized using a stacking metaphor: the most proximate window is the top or focus window, which contains the active document that the user is working on. Since in PaperTab, windows are laid out on a physical desk with more real estate, stacking of windows need not be the dominant metaphor for task focus. Rather, and according to Sellen and Harper's analysis of paper document use, the focus or activity level of a window is determined by the proximity of a tab to the user [30]. Figure 3 shows PaperTab's three zones of proximity to the user, each pertaining to a different focus level: *hot* (within arm's reach, active document), *warm* (at arm's reach, locked or stacked document), and *cold* (outside arms reach, filed document).

Hot Zone

In this zone, tabs are either held by the user, or within immediate hand's reach of the user. They are the equivalent of the top window in GUIs. They contain active documents editable by the user via touch input or keyboard. When a tab is moved into the hot zone a small LED in the top left corner of the display turns green indicating they have become editable. When displays are moved to the hot zone is there no change to the view resolution. This is only true for documents moved to the hot zone, and the reason is to allow users to examine the contents of warm and cold tabs without changing their view. When tabs are moved in the other direction, i.e., from the hot zone to the warm zone, their view changes to a thumbnail overview of the document. When they are moved to a cold zone, the document closes to show a file icon. This is the equivalent of closing a window in a GUI. Tabs remain hot until the user releases them in a warm or cold zone. Users can use touch or bend navigation to change view in the hot zone. E.g., in the hot zone, users can go from a thumbnail overview to a full screen view by bending the sides of the display outwards.

Warm Zone

In this zone, tabs are at arm's reach of the user. They are the equivalent of minimized windows, or windows stacked below the top window in a GUI. As such, documents contained in warm tabs are locked, and not directly editable by the user. This allows them to be handled and placed in piles without the risk of spurious input. They do, however, respond to external events such as alerts, incoming files, emails, or edits by remote users. Whenever a tab is moved to the warm zone this is indicated by an LED in the top left corner of the display turning amber. When a tab is moved from a cold zone to the warm zone, the documents in the cold tab are opened into tabs displaying thumbnail overviews. This allows a convenient way of opening up multiple documents onto a single tab. These tabs can be selected through touch when the tab is picked up. When the tab is moved back to the cold zone, the tabs close to show a file icon. When a hot tab is moved to the warm zone, this causes the tab to show the thumbnail overview of its document. For example, if the user is browsing a full screen photo in an album, moving the display into the



Figure 4. Pointing with a hot display (right) into a warm display, in this case showing a TOC (left), displays detailed focus information on the hot display, in this case a book page.

warm zone would cause it to display thumbnails of all photos in the album.

Cold Zone

In this zone, tabs are just outside of arm’s reach of the user, yet easily accessible by leaning or reaching forward over the desk. Cold tabs allow storage of documents out of the way of the active task. They are equivalent to *file folders* in GUI filing systems. The cold zone provides users with an easy way to file and retrieve documents. Filed folders are not editable and only respond to touches by other tabs. When a tab is moved into the cold zone, the LED in the top left corner of the display turns off, but it continues to respond to touch input. The display of a tab in the cold zone is powered off when not interacted with and, similar to printed documents, does not consume energy. Displays are powered back on when they are picked up, touched by another tab, or updated externally.

Document Notifications

When a document on a warm tab receives an external update, its LED starts flashing amber to notify the user. E.g., a user may keep her E-mail Inbox in a warm tab displaying a list of recent emails. When a new email arrives, this tab starts blinking. Users can open the last received email by moving the tab with the Inbox document into the hot zone and touching the email. Moving the Inbox back into the warm zone closes the current email and moves the window back to its original list overview. The LED notification stops blinking when the tab is touched.

Focus+Context Interactions

Multi-display interactions provide convenient focus+context navigation of content on hot, warm, or cold tabs via a second, hot tab. When a hot tab is moved into another zone during this interaction, it remains hot as long as the user holds the tab. When a user points the top left corner of the hot tab onto the thumbnail on the underlying warm or cold tab, the hot tab shows a full screen preview of the thumbnail. This is the

equivalent of magic lenses [28], hyperlink previews, or contextual menus in GUIs. After previewing the item, the user can move it permanently onto the hot tab by lifting it and pulling it into the hot zone [28], preserving the detailed view. Users can also point within hot tabs. E.g., pointing at a URL in a document on one hot tab shows a full-screen preview of the webpage on the top hot tab. Pointing at a location on a map or an item in a book’s table of contents may pop up a linked page on the second screen (see Figure 4). This technique provides an interesting alternative to the Pick and Drop technique [26], for multiple displays.

Filing and Opening Documents

Document files are opened by pointing an empty tab at a list of document file icons represented on a cold zone folder. During pointing, the hot tab shows a thumbnail overview of the contents of the file document. In the cold zone, file icons can be moved between folders by pointing at their thumbnail, picking them up onto the hot tab, and then tapping a third tab. When a user points with a hot tab at an empty space within a file folder in the cold zone, its document is closed and the thumbnail is moved into the filing folder. This closes the tab, emptying its screen.

Moving Data Objects

Users can also copy or move documents and data objects within documents via this technique. This action is equivalent to using a GUI clipboard for cutting, copying and pasting files. For example, this way users can add an attachment to an email displayed on one tab by tapping its empty space with a pdf document on another tab. Alternatively, users can move a data object between two tabs by placing them adjacent to each other, then dragging the data object from one display to the other with a finger.

Keyboard Input

Users can type on a standard wireless Bluetooth keyboard to edit text. The location of the keyboard on the desk is tracked, and input is automatically routed to the document that is closest to the keyboard. Users can tie the keyboard to a tab by hitting a function key on the keyboard while the keyboard is adjacent to it. When users subsequently move the tab, input continues to be directed to it.

Bend Interactions

Individual tabs respond to bends on the top left and right corners of the display, as well as the full display. According to Lahey et al. [18] users have a preference for top corner bends over bottom corner bends. Building upon these results, we implemented the following bend gestures for PaperTab: Bi-directional top-right corner bends for navigation, and bi-directional top-left corner bends for application-specific tasks (eg., the tear gesture to unlink two collocated tabs, or to reply to an email). Full display bends are used for zooming in and zooming out.




Figure 5. Exploded view of a tab, which is made out of 6 layers: 1) Flexible Touch Layer, 2) Flexible Display, 3) 6 DOF Tracker, 4) Bend Sensor Layer and 5) Capacitive Grounding wire (not shown).

A downward top-right corner bend selects the next tabbed document or pages forward, while an upward top-right corner bend pages back [35]. Bends are also used to scroll through lists that are larger than fit the physical display. Bends performed across two tabs copy the selected information from the top to the bottom tab. Users can stack tabs and bend the top right corners inwards with a single gesture to create a file folder containing the contents of each of the tabs: the folder appears on the top tab and the other tabs become empty.

Colocation of Tabs

If an empty tab is placed directly adjacent to and slightly overlapping with a tab, the system responds by creating a larger view of the original document across the two displays [14,16]. Items can be moved or copied between two collocated tabs by dragging them via touch. If collocated tabs are moved away from each other, they will display a subset of the larger graphics environment that is framed by the two displays [9]. This allows browsing of large graphics documents using multiple tabs as magic lenses. E.g., if one tab containing a Google map of a city is collocated with another, the map expands to encompass the area next to it. If the second tab is now moved away, the view of that tab will automatically scroll to the relative distance from the original tab, thus revealing other cities. This can be useful when planning detailed routes between an origin and destination city. This feature is also useful for browsing large graphical documents, such as architectural drawings, across the entire desk space without zooming. Collocated tabs are disconnected again using a tear gesture. This consists of bending the top left corner of right tab upwards while moving the display upwards. Upon tearing, two separate documents are created, each containing the data objects on display. In the case of the map example, this would, e.g., produce one tab with the origin city, and one with the destination city.

Summary of Interaction Techniques

To summarize, PaperTab’s interaction techniques include the following key elements:

Zones: Tabs are either hot (editable), warm, or cold (off) depending on their distance to the user.

Document-centrism: A tab may display no documents, part of a document, a single document, or multiple documents (tabbed).

Varying detail: The level of detail displayed by tabs changes when moved between zones. The only exception is when tabs are moved directly to the hot zone, in which case there is no change of view.

Focus+Context: The tip of a tab can be pointed at the contents of other tab to show more detailed views of data objects or files. Tabs can also be collocated to create larger views.

Input: Interactions with tabs occur through touch, keyboard, bending, proximity (to other tabs, to the desk, and to the user), as well as by pointing between two tabs.

PAPERTAB IMPLEMENTATION

Figure 1 shows PaperTab, consisting of a desk, delineated into virtual hot, warm, and cold zones. While there is no real limit to the number of tabs that can be placed on the desk, our current prototype, as shown, features 9 tabs.

Tab Architecture

Figure 5 shows an exploded view of a tab consisting of 4 flexible layers. The top layer consists of a flexible capacitive touchscreen. The second layer contains the flexible electrophoretic display. The third sensor layer provides tracking, while the fourth senses bending of the corners of the display and also consists of a circuit that allows capacitive coupling between tabs for pointing. Each tab is approximately .5 mm thick and weighs approximately 20 g. Each tab is a dumb

terminal that communicates with a host PC for UI updates, sensor input and to coordinate data exchanges with other tabs.

Layer 1: Flexible Touch Input

The first layer consists of a flexible, thin-film, transparent Zytronic capacitive touch screen [36], 10" in size diagonally. This layer is connected to a host PC with a ribbon cable, via which it issues touch events to the application software running on the tab. In order to enable multi-display pointing using the touch sensor, we developed a capacitive circuit that, when the tab is held by the hand, transfers the user's capacitance to a soldered tip on the top left corner of the tab. This activates touch events upon one tab touching another.

Layer 2: Flexible Electrophoretic Display

The second layer features a large 10.7" diagonal flexible Plastic Logic electrophoretic display with a minimal bezel that was custom manufactured for the PaperTab system. The display features a resolution of 1280 x 960 pixels, with a full-screen refresh rate of 700 ms. The display is connected to a custom PlasticLogic driver board, underneath the desk, via ribbon cables. This board is connected over USB to a PC that controls the system logic and the user interface.

Layer 3: 6 DOF Location and Orientation Tracking

An electromagnetic sensor mounted on each tab allows tracking of location and orientation relative to the other tabs, as well as the desk and the user. It consists of a small trakSTAR sensor probe that is attached via a wire to a processor box placed underneath the desk [1]. While displays could easily be tracked wirelessly via computer vision systems available to the authors (e.g., Vicon, Kinect) [16], we chose to track them using a wired configuration in the current prototype because of the following reasons. Firstly, computer vision systems suffer from occlusion, making it difficult to track tabs when they overlap. Secondly, we believe an electromagnetic solution to provide a better value proposition for production models. Finally, the electromagnetic tracker provides greater resolution in all 6 DOF.

Hot, Warm and Cold Zone Tracking

Our software translates the 6 DOF coordinates of each tab to determine the location of a tab relative to the user, relative to the desk, and relative to other tabs. A tab is hot when it is lifted from the table, or when it is within arms length (70 cm) from the user. A tab is warm when it is on the desk between 70 cm and 110 cm from the user. All other tabs are cold, with their displays powered off. These measurements are based on the affordances of the desk and size of the displays. They can be easily reconfigured to suit different desk sizes and usage patterns. The electromagnetic tracker also determines when tabs are collocated, and tracks the distance between them when connected tabs are moved apart. Finally, the electromagnetic tracker determines when tabs are stacked on the basis of their z coordinate.

Layer 4: Bend Sensor Layer

The bend sensitive layer consists of 2 bi-directional FlexPoint [10] bend sensors mounted on a custom-built

flexible circuit board mounted directly underneath the display. Bend sensors are connected via a flat ribbon cable to an Arduino Mega 2560 prototyping board that communicates with the host PC for processing bends.

Layer 5: Capacitive Grounding Wire (not shown)

The bottom layer features only a small wire connected to a small metal tip that transfers the capacitance of the user's hand to another displaywindow for the purpose of activating the Zytronic touch film. This is used for focus+context pointing with one displaywindow onto another.

PaperTab System Setup

Tabs are tethered to a host PC, an 8-Core MacPro running Windows 7 that is placed underneath the desk. A C# host application runs all the interaction logic, manages the system state and pushes UI updates to the tabs. Bend sensor input and touch sensor data are processed by the C# application. A trakSTAR processor/emitter mounted underneath the desk preprocesses the 6 DOF sensor data and send coordinates back to the C# host. The processed data from all these sensors are used for determining the state of the system and tracking interactions in real-time. The user interface for each tab is rendered on the host PC's buffer and pushed out to each tab via the driver boards.

USER STUDY

We conducted a qualitative user study to elicit feedback on PaperTab. Twelve participants (5 female, aged 21-30, all right-handed) participated in think-aloud sessions lasting an hour. Each session began with a 5-minute introduction followed by a 10-minute period during which users explored the system. Participants were then given instruction on any functionality that they did not discover. Finally, each participant was asked to perform 12 common computer usage tasks using PaperTab. After the tasks, users were asked to rate different interaction techniques for Learnability, Efficiency, Physical Demand and Mental Demand on a 5-point Likert scale (Strongly Disagree-Strongly Agree). Their general observations were noted.

Tasks

We first asked each participant to place the tabs on the desk, based upon when they should expect to next interact directly with the display: (1) in a moment (2) within the next few days, and (3) in the next few weeks. The primary purpose for this task is to understand users' preferred locations for the hot/warm/cold zones. We recorded the positions of the tabs and used them to calibrate the location of the various zones on the desk.

Then, participants were asked to work on twelve different tasks, which were designed to increase the need for a group of devised interactions with multiple tabs. The twelve tasks included:

a) *Document management*: Participants interacted with several documents by pointing, bending, moving, and col-

Interaction Techniques	Easy to Learn	Efficient	Physically Demanding	Mentally Demanding
Changing the display view between zones	4.25 (0.22)	3.92 (0.26)	2.5 (0.44)	2 (0.35)
Point with display at another display	4.33 (0.19)	4.42 (0.15)	1.5 (0.26)	1.5 (0.26)
Bending the side or corner of a display for navigation	4.25 (0.22)	4.33 (0.19)	1.92 (0.26)	2.08 (0.34)
Overlapping two displays to extend view	4.33 (0.26)	4.25 (0.22)	1.83 (0.30)	1.42 (0.15)
Overall impressions	3.92 (0.31)	4.08 (0.19)	2 (0.30)	2.25 (0.35)

Table 1. Mean ratings and standard errors (s.e.) for responses on a selection of interaction techniques. (1 - Strongly Disagree, 5 - Strongly Agree)

locating tabs on the desk. They were asked to perform four tasks: (1) Search for a described artifact (an image or a phrase) from one of three documents; (2) Navigate to a specific page in a document; (3) Find total number of artifacts in a given document; (4) Compare two pages of a document based on a given criteria.

c) Answering Email: Participants used pointing, bending, navigation, and pick and drop interactions. They were asked to perform 3 tasks: (1) find an email with a specified attachment; (2) find requested information from one of the emails in the inbox; and (3) reply to an email by attaching a specified photo.

d) Route Planning: Participants used pointing, bending, collocating, and panning interactions to plan a route using a map. They were asked to perform 3 tasks: (1) Find a location on the map; (2) Collocate two tabs to extend the map view; (3) Pan to a specified location on extended tab.

RESULTS

Participants learned to perform the interactions in less than 5 minutes. Table 1 shows the results from our questionnaires. A χ^2 test was used to evaluate responses. To ensure an estimated frequency of 5 per cell we collapsed data into 2 categories: positive and negative, ignoring neutral responses. Overall, users found PaperTab easy to learn ($\chi^2(1)=7.364$, $p=0.007$) with ratings in the order of 4 out of 5, efficient, with ratings in the order of 4 out of 5 and less physically demanding with rating 2 out of 5 ($\chi^2(1)=5.444$, $p=0.020$).

Moving between Zones

Changing the display view between zones was found to be easy to learn, with ratings in the order of 4 out of 5 and efficient, with ratings in the order of 4 out of 5. When asked to place tabs on the desk, based upon when they would be used again, for hot and warm zones, participants placed the tabs in locations consistent with our zone interactions. However,

5 participants noted that they would prefer to keep the displays away from the desk for cold zone documents. One participant proposed stacking a display that would not need to be accessed for several days beneath another display located in the warm zone.

Rather than using bends, one participant proposed having an “Outbox zone” to place emails for sending. Other comments included: “I liked that when I was pulling it closer it knew I wanted to read in more detail and pulled up the document” and “Once aware of zones, I was at first concerned about moving between them, but I realized that no data is lost and the display changes back when I pick it up again.” Another participant wanted to be able to lock or unlock a display, by bending a corner of the display, to allow for more control over the zone behavior.

Pointing within Tabs

Users found Pointing with a display at another display to be easy to learn, with ratings in the order of 4 out of 5, efficient, with ratings in the order of 4 out of 5, less physically demanding, with ratings in the order of 1 out of 5 ($\chi^2(1)=8.333$, $p=0.004$), and less mentally demanding, with ratings in the order of 1 out of 5 ($\chi^2(1)=8.333$, $p=0.004$). 3 Participants mentioned in their open-ended questionnaire responses that they found pointing to be superior to dragging: “Why should I have to drag when I can point?”. Dragging involves friction and often required the user to bend over the display. When prompted for further elaboration, participants noted that pointing with a tab was easier, as they did not have to worry about accidentally releasing a drag operation in the wrong location (despite the fact that objects would snap back upon erroneous release) and generally required less effort. We observed that pointing with a display appeared to require less focus and fewer fine motor movements than dragging.

Most participants were able to understand how to perform the interactions with minimal prompting. For example, after

explaining that tabs could be used to point at each other, participants were able to determine, without any further instruction, how to attach documents and also put them back into folders. Two participants wrote that they particularly enjoyed this interaction: “*I very much enjoyed using the tap function to attach documents*”. However, we also observed that pointing with the edge of a tab could be challenging depending on the location of, and grasp on, the displays. The tab used as a pointing device sometimes occluded the other tab, requiring the participant to reposition their hand. One participant suggests that the right corner should also be able to point – less movement would often be required if this functionality was added.

Other Interactions and General Feedback

Bend interactions were found to be easy to learn, with ratings in the order of 4 out of 5, efficient, with ratings in the order of 4 out of 5, and less physically demanding, with ratings in the order of 2 out of 5. Participants wanted to be able to change the polarity of interactions (bend up to page forward vs. bend down to page forward). 2 participants indicated that they did not prefer to bend the top-left corner and wanted to have most bend interactions confined to the right side of the display.

Overlapping display to extend it was found to be easy to learn, with ratings in the order of 4 out of 5 ($\chi^2(1)=8.333$, $p=0.004$), efficient, with ratings in the order of 4 out of 5 less physically demanding ($\chi^2(1)=8.333$, $p=0.004$), with ratings in the order of 1 out of 5 ($\chi^2(1)=8.333$, $p=0.004$), and less mentally demanding, with ratings in the order of 1 out of 5 ($\chi^2(1)=6.400$, $p=0.011$). 4 participants cited collocation as a feature they would use often. Participants mentioned that, unlike standard GUIs, enlarging by co-location does not come at the expense of other display content.

DISCUSSION AND FUTURE DIRECTIONS

Overall, our results were in line with expectations. PaperTab interaction techniques appear easy to learn and do not induce significant physical demands, while achieving low ratings for mental load. One of the most common suggestions from participants trying our prototype is that it should make more extensive use of zones for different document types. For example, a video could be played full screen in the hot zone, paused in the warm zone, and appear as a thumbnail with a description in the cold zone. However, at the same time, users appreciated the one-directionality of view changes for hot tabs. It appears that automated view changes based on proximity are a balancing act of allowing easy overview of information while keeping the user in control of the view. Feedback from the user study indicated strong preference for pointing with tabs over dragging via touch. We may want to improve the surface characteristics of the touch film in future versions.

Although the tabs are thin and easy to manipulate, our current prototype tabs still suffer from wired tethering. While tethering of tabs appeared awkward to users, it generally did not seem to interfere with interacting with the system. Future

versions will remove the need for tethering by introducing lightweight processors attached on the side of each tab (as a rigid handle), and by processing of touch, bend and tracking operations on the tab itself.

CONCLUSIONS

In summary, we presented PaperTab, a paper computer with 9 letter-sized touch-sensitive, flexible, electrophoretic displays. The displays, augmented with sensors and driver boards, are called tabs. Each tab embodies a single document window. The location of each tab is tracked on a desk using an electromagnetic tracker. Document views are adjusted when the tab is moved between three proximity-based zones. Tabs can be pointed at one another for focus+context operations, while touch and bend sensors allow users to navigate content. User study results suggest that our interaction techniques are overall easy to learn and less physically demanding.

REFERENCES

1. Ascension, Inc. trakSTAR. <http://www.ascension-tech.com/medical/trakSTAR.php>
2. Agrawala, A., and R. Balakrishnan. Keepin' it Real: Pushing the Desktop Metaphor with Physics, Piles and the Pen. In *Proc. CHI'06*, ACM Press (2006), 1283-1292.
3. Baudisch, P., Good, N., Bellotti, V., and Schraedley, P. Keeping things in context: a comparative evaluation of focus plus context screens, overviews, and zooming. In *Proc. CHI*, ACM (2002), 259–266.
4. Bier, E. a, Stone, M.C., Pier, K., Buxton, W., and DeRose, T.D. Toolglass and magic lenses: the see-through interface. *Proc. SIGGRAPH*, (1993), 73–80.
5. Bush, V. As We May Think. *The Atlantic Monthly* 176, 1 (1945), 101-108.
6. Cockburn, A. and McKenzie, B. Evaluating the effectiveness of spatial memory in 2D and 3D physical and virtual environments. *Proc. CHI*, (2002), 203.
7. Cockburn, A. Revisiting 2D vs 3D implications on spatial memory. *Proc. AUIC*, (2004), 25–31.
8. Fitts, P.M. The information capacity of the human motor system in controlling the amplitude of movement. *J Exp Psychol* 47, 6 (1954), 381-391.
9. Fitzmaurice, G.W. Situated information spaces and spatially aware palmtop computers. *Communications of the ACM* 36, 7 (1993), 39-49.
10. FlexPoint Inc. <http://www.flexpoint.com>
11. Girouard, A., Tarun, A., and Vertegaal, R. Display-Stacks: interaction techniques for stacks of flexible thin-film displays. In *Proc CHI'12*. (2012). ACM, New York, NY, 2431-2440.
12. Greenberg, S., Marquardt, N., Ballendat, T., Diaz-Marino, R., and Wang, M. Proxemic interactions: the new ubicomp? interactions 18, 1 (2011), 42-50.
13. Guimbretiere, F. Paper augmented digital documents. *Proc. UIST* 5, 2 (2003), 51-60.

14. Hinckley, K., Ramos, G., Guimbretiere, F., Baudisch, P., and Smith, M. Stitching: pen gestures that span multiple displays. *Proc. AVI*, (2004), 23–31.
15. Hinckley, K., Yatani, K., Pahud, M., et al. Pen + Touch = New Tools. *Proc. UIST*, (2010), 27-36.
16. Holman, D., Vertegaal, R., Altosaar, M., Troje, N., and Johns, D. PaperWindows : Interaction Techniques for Digital Paper. *Proc. CHI*, (2005), 591-599.
17. Khalilbeigi, M., Steimle, J., and Mühlhäuser, M. Interaction techniques for hybrid piles of documents on interactive tabletops. *Proc. CHI extended abstracts*, (2010), 3943.
18. 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).
19. Lee, J.C., Hudson, S.E., and Tse, E. Foldable interactive displays. *Proc. UIST*, (2008), 287.
20. Luff, P., Heath, C., Norrie, M., Signer, B., and Herdman, P. Only touching the surface: creating affinities between digital content and paper. *Proc. CSCW*, (2004), 523.
21. Mackay, W.E., Pothier, G., Letondal, C., Bøegh, K., and Sørensen, H.E. The missing link: Augmenting Biology Laboratory Notebooks. *Proc. UIST*, (2002).
22. Merrill, D., Kalanithi, J., and Maes, P. Siftables: towards sensor network user interfaces. *Proc. TEI*, (2007), 75–78.
23. Mine, M.R., Brooks, F.P., and Sequin, C.H. Moving objects in space: exploiting proprioception in virtual-environment interaction. *Proc. SIGGRAPH*, (1997), 19-26.
24. Plastic Logic Inc. (2012). <http://www.plasticlogic.com>
25. Reilly, D., Rodgers, M., Argue, R., Nunes, M., and Inkpen, K. Marked-up maps: combining paper maps and electronic information resources. *Personal and Ubiquitous Computing* 10, 4 (2005), 215-226.
26. Rekimoto, J. Pick-and-Drop : A Direct Manipulation Technique for Multiple Computer Environments. *Proc. UIST*, (1997), 31-39.
27. Rekimoto, J., Ullmer, B., and Oba, H. DataTiles. *Proc. CHI*, (2001), 269-276.
28. Rohs, M., Oulasvirta, A., and Suomalainen, T. Interaction with magic lenses: real-world validation of a Fitts' Law model. *Proc. CHI*, (2011), 2725–2728.
29. Schwesig, C., Poupyrev, I., and Mori, E. Gummi: a bendable computer. *Proc. CHI*, (2004), 263 - 270.
30. Sellen, A.J. and Harper, R.H.R. *The Myth of the Paperless Office*. MIT Press, 2003.
31. Terrenghi, L., Kirk, D., Sellen, A.J., and Izadi, S. Affordances for manipulation of physical versus digital media on interactive surfaces. *Proc. CHI'07*, (2007), 1157–1166.
32. Vertegaal, R. and I. Poupyrev. Organic User Interfaces: Introduction to Special Issue *Comm of the ACM* 51(6), (2008), 26-30.
33. Ullmer, B., Ishii, H., and Jacob, R.J.K. Token+constraint systems for tangible interaction with digital information. *TOCHI* 12, 1 (2005), 81-118.
34. Wellner, P. Interacting with paper on the DigitalDesk. *Communications of the ACM* 36, 7 (1993), 87–96.
35. Wightman, D., Ginn, T., and Vertegaal, R. BendFlip: Examining Input Techniques for Electronic Book Readers with Flexible Form Factors. *Proc. INTERACT*, (2011).
36. Zytronic, PLC . (2012). <http://www.zytronic.co.uk>