Piles, Tabs and Overlaps in Navigation among Documents

Mikkel Rønne Jakobsen and Kasper Hornbæk
Department of Computer Science, University of Copenhagen
Njalsgade 128, Building 24, 5th floor, DK-2300 Copenhagen, Denmark
{mikkelrj, kash}@diku.dk

ABSTRACT
Navigation among documents is a frequent, but ill supported activity. Overlapping or tabbed documents are widespread, but they offer limited visibility of their content. We explore variations on navigation support: arranging documents with tabs, as overlapping windows, and in piles. In an experiment we compared 11 participants' navigation with these variations and found strong task effects. Overall, overlapping windows were preferred and their structured layout worked well with some tasks. Surprisingly, tabbed documents were efficient in tasks requiring simply finding a document. Piled documents worked well for tasks that involved visual features of the documents, but the utility of recency or stable ordering of documents was task dependent. Based on the results, we discuss the effects of spatial arrangement, visibility, and task-dependency, and suggest areas for future research on document navigation and its support by piling.

Author Keywords
Document navigation, window switching, overlapping windows, piled windows.

ACM Classification Keywords
H.5.2 User Interfaces (D.2.2, H.1.2, I.3.6): Graphical user interfaces (GUI).

INTRODUCTION
Users frequently navigate between multiple documents, for instance to copy-paste text, compare web pages, or capture notes when reading. But navigation between documents is hard. Multiple windows or tabs are increasingly used in web browsing, for example, but users find it hard to manage multiple web pages [31]. In programming, source files are often viewed one at a time, requiring frequent navigation between files. Tabbed document interfaces that are common in widespread programming environments, such as Eclipse and Visual Studio, make navigating among many source files cumbersome [21].

NordiCHI 2010, October 16–20, 2010, Reykjavik, Iceland.
 ISBN: 978-1-60558-934-3...$5.00.

Much research has aimed at supporting general window switching and switching between tasks or activities. General window switching techniques [26,29] help switching between applications, which are often visually distinct. However, switching between documents often takes place within one application and documents switched between may be visually similar. Also, while users' frequent switching between tasks or activities needs support [7,18,27,28,30], often users must switch between multiple documents related to one task.

Problems with existing switching mechanisms
One facility for supporting navigation among documents is overlapping windows, standard in most operating systems. Hutchings et al. [19] found that 78% of the time people had eight or more windows open, making locating a particular window within a group of overlapping windows difficult. Switching between windows using the taskbar is also hard because users may only see a short part of the windows' titles and because the taskbar occasionally collapses documents into one tile [13].

Another facility, using tabs to represent multiple documents within a single window, is seen in many web browsers, spreadsheets, and programming environments. But tabbed interfaces can only show a limited number of documents at any time, necessitating further interaction with scrollbars or drop-down lists. When several documents are open, tabs provide minimal visual cues for finding a particular

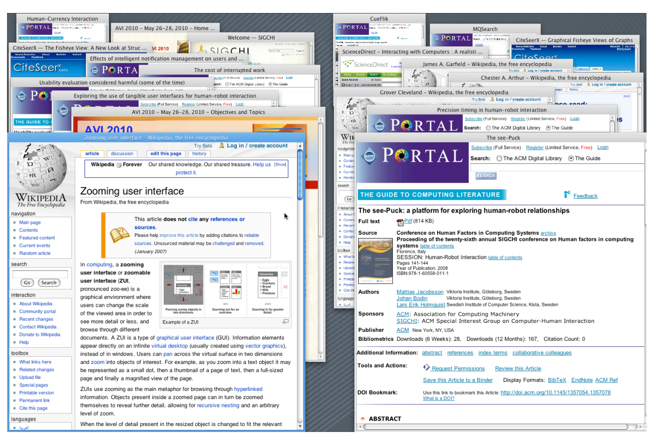

Figure 1: Interface with two piles of spatially arranged documents.

document. Although recent studies show increased use of tabs in web browsing [15,31], attempts at improving or reinventing tabbed interfaces are needed [3]. Beaudouin-Lafon [6] suggested techniques for interacting with tabbed documents, such as leafing through tabs, but these techniques have to our knowledge not been widely adopted.

Still other facilities support navigation among multiple documents by a notion of piles. People pile and spatially arrange paper documents to see many of them at once, to structure their tasks, or to remind [22,25]. Although techniques for working with piles of electronic documents have been explored [4,5,6,23], questions about the design of interfaces that use piling remain unanswered. For example, how can a pile be spatially arranged so as to support navigation between its documents? Furthermore, empirical studies that investigate the usefulness of piling electronic documents are lacking.

This paper explores the design space of interfaces that support document piling and investigates possible benefits of spatially arranged piles (see Figure 1). In an experiment, we compare variations of interfaces for supporting navigation between documents, including two that arrange documents in a pile. The results contribute insights about the effect of visibility and spatial arrangement in different tasks involving navigation between documents. These findings may inform design of techniques for navigating between documents and indicate areas where further empirical research is needed.

SUPPORT FOR DOCUMENT NAVIGATION

Studies of office work have found that people benefit from piling as a way of spatially grouping documents [22,23,25]. Arranging and piling electronic documents may also help users navigate among documents. However, methods for spatially arranging documents and techniques for interacting effortlessly with them are needed.

The use of electronic piles has been researched mainly as an alternative to filing documents by storing them as named files in hierarchical folders, and work on iconic document representations [4,5,23]. Instead, we focus on using piles for supporting navigation among documents that are in use.

Based on a review of the literature, we have developed a taxonomy that incorporates important dimensions in the design of document piling interfaces (see Table 1). We relate each dimension to previous work and we explain the approach taken in the design of the piling interfaces used in the experiment presented later in the paper.

Spatial arrangement

Documents can be spatially arranged to help users navigate among them. Four concerns of spatial arrangement in design of piling interfaces are visibility, spatial memory, scanning, and focus/periphery.

Visibility of documents reminds users of the documents and allows them to interact with the documents. In overlapping window interfaces, windows open in a cascade to keep from

Documents	spatially arranged for	**visibility** **spatial memory** scanning focus/periphery	
	arranged	manually **automatically**	to increase user control to reduce user efforts
	viewed	**permanently** transiently	to provide visual cues to reduce use of space
	represented	full-size with content **scaled**/cropped	for content interaction for overview
	arranged in	**fixed** **adaptive**	order

Table 1: Taxonomy of document interface designs that use piling. The terms describing the interfaces used in the experiment are highlighted in bold.

hiding the title bar of windows underneath, as illustrated in Figure 2 (a).

Spatial memory can be utilized to help people organize and find documents. Figure 2 (b) illustrates a spatial arrangement of documents in a pile that makes the spatial location of individual documents more distinct. Data Mountain, which allows users to arrange web page thumbnails on a perspective plane, showed faster and more accurate retrieval of pages compared with a text-based bookmark list [14].

Scanning through documents in a serial manner may be easier if the documents are aligned. Compared with the arrangement in Figure 2 (b), the alignment in Figure 2 (a) and (c) may provide for easier scanning of the documents. Agarawala et al. explore pile interaction techniques that use various aligned layouts for leafing through piles in a serial manner [4].

(a) (b) (c)

Figure 2: Spatial arrangements of four documents. (a) Pile of documents with titles and part of content visible, aligned diagonally. (b) Pile of documents with titles and part of content visible, arranged irregularly. (c) Document window with titles partly shown in tabs; tabs aligned horizontally.

Focal/peripheral areas seem to play a role in the spatial layout of documents [25] and how users manage space in large or multiple displays [8,16]. Robertson et al. divide the display into focal and peripheral regions in a window interface that supports task switching. Their Scalable Fabric uses animated transitions to shrink windows to thumbnails as they are moved into the peripheral region [27].

In this paper, we arrange a pile of documents vertically so as to make the upper part of each document in the pile visible, including its title. To further improve visibility in piles with many documents, we reduce overlap by staggering the documents, which may also help users find documents by spatial location.

Manual vs. automatic layout

Document interfaces must provide the user with control of the layout to arrange the documents in ways that are appropriate for different tasks, with as little effort as possible.

Manual layout in overlapping window interfaces allows users to pile windows manually—users have full control over moving, resizing, and overlapping windows. However, arranging windows and switching between overlapping windows require time-consuming window management [10]. Compared with overlapping windows, automatically tiled windows may increase performance, but users may prefer overlapping windows [9,20].

Automatic layout of documents in piles may help users arrange documents and navigate among them with fewer efforts. Task Gallery provides an “ordered stack”, in which windows are automatically laid out with fixed distance between them, as well as a “loose stack”, which users lay out manually [28]. Clicking on a window in one of these stacks moves the window forward to a selected window position. The study focused on the system’s task management support and it is not clear how the stacks were used. Lightweight interactions have been explored for automatically gathering a selection of items into a pile on the desktop, for manipulating piles, and for arranging piles in different layouts with few user efforts [4]. However, it is not clear how these techniques would work for piles of document windows.

In this paper, we explore piling interfaces that automatically maintain the layout so that all documents in a pile, and in particular the documents’ titles, are visible. When the user opens or closes a document, or drags a document from one pile to another, the piled documents rearrange to maintain a consistent layout of documents. We thus aim to reduce the user efforts required for rearranging and navigating among documents in piles compared with overlapping windows.

Permanent views vs. transient views

Permanent views that support navigation among documents provide visual cues of the documents. For instance, in many desktop environments the task bar often shows windows as icons or text, including windows that are hidden.

Transient views that can be called up temporarily can support navigation without permanently using display space. Widespread examples of transient views are the Exposé feature on Mac OS X [1] and the <alt>-<tab> window switching technique. Several techniques have been explored for browsing and leafing through piles that use transient views of the piled documents [4,5,23]. However, we are unaware of studies that compare how users navigate among piled documents with permanent and transient views.

In this paper, we arrange documents in a pile to provide permanent visual cues for navigation so as to compare piles to permanent views techniques for navigation between documents that are common in widespread interfaces.

Document representation

Full-size views of documents allow the user to interact directly with the documents’ content. In contrast, small document representations using thumbnail views, icons, or text make many documents visible in limited display space.

Windows can be reduced for instance by scaling the windows [11,27,28], or by scaling or cropping the content of windows [24]. Using a 3D document space, The Web Forager showed a document at a focus position in full size for direct interaction between user and content, whereas documents that are not the immediate focus are placed at greater z-distances (and thus smaller) [11]. Aiming to reduce screen clutter, Miah and Alty studied an adaptive window management system that reduced the size of windows [24]. They found that identification of a window became difficult when the window’s size was below 30% screen utilization. Their results also indicate marginally higher accuracy in identification tasks when a window’s content is cropped rather than scaled. Users may identify a document by a thumbnail view from among visually distinct documents (e.g., pages from different web sites). However, if documents are visually similar (e.g., source code files or pages from digital library), a thumbnail view contains no salient features for identifying the document.

Semantic zooming or use of colors can make visually similar document thumbnails distinguishable. Other features of the document representation can help navigating among documents in piles. For instance, to help finding one among many equally-sized windows piled on top of each other, Beaudouin-Lafon suggests rotating windows [6].

In this paper, we change the size of documents in a pile so that documents at the bottom of the pile are smaller. To reduce a document’s size, we scale the document’s content, but show its title in a fixed font size so that the title remains readable. The readable title is aimed at helping users in searching among visually similar documents.

Adaptive vs. fixed ordering

The order that documents are shown in the display may influence how users navigate between the documents.

In most desktop environments, windows overlap each other in the order they were last recently used—and windows are shown in this order when switching using <alt>-<tab>. Recent research adapts the order of windows used in window switching techniques based on semantic or temporal information about the windows [26,29]. These examples use *adaptive ordering*: the order of windows changes as the user switches between the windows.

Alternatively, documents can have a *fixed order*. In this case, the user may order documents manually or documents may be ordered automatically by the system (e.g., by creation date).

In this study we investigate different ways of ordering documents in a pile to see how that influences the way users navigate among the documents. We use a fixed ordering where documents are piled in the order they were opened and an adaptive ordering where documents are piled in the order they were most recently used.

Summary

We have described the design space of interfaces for document piling based on five dimensions in the taxonomy shown in Table 1. Using this taxonomy, we have explained the approach taken in the interfaces used in the experiment. Earlier studies of document piling have mainly focused on piling document files in desktop environments, use of iconic representations, and use of transient view techniques to provide visibility of a pile's documents. In contrast, we focus on piling documents in use, documents represented as zooming windows, and spatial arrangement to provide visibility of a pile's documents. We are unaware of earlier studies of automatically arranging piles of documents that aim at providing permanent visual and spatial cues for navigating among the documents.

EXPERIMENT

To investigate the influence of spatial arrangement and visibility of documents on how users navigate among documents, we conducted a controlled experiment in which four interfaces for viewing multiple web pages were compared. We use web pages because the Web is a widely used source of information, for instance in sense-making of research literature using digital libraries [32], and because users often have many web pages open simultaneously [31].

Interfaces

The interfaces consist of windows, where each window contains only the view of a web page, framed by a border and a title bar at the top. The active window is indicated by blue. Windows have a fixed size of 800 x 700 pixels.

The Overlapping interface

In the Overlapping interface (shown in Figure 3), windows can be moved by dragging the title bar. Document windows open in a cascade to avoid hiding title bars of windows underneath. A horizontal bar (see below) at the bottom of the display contains a tile for each document (similar to the Taskbar in Windows XP or Vista [2]). Clicking on a tile

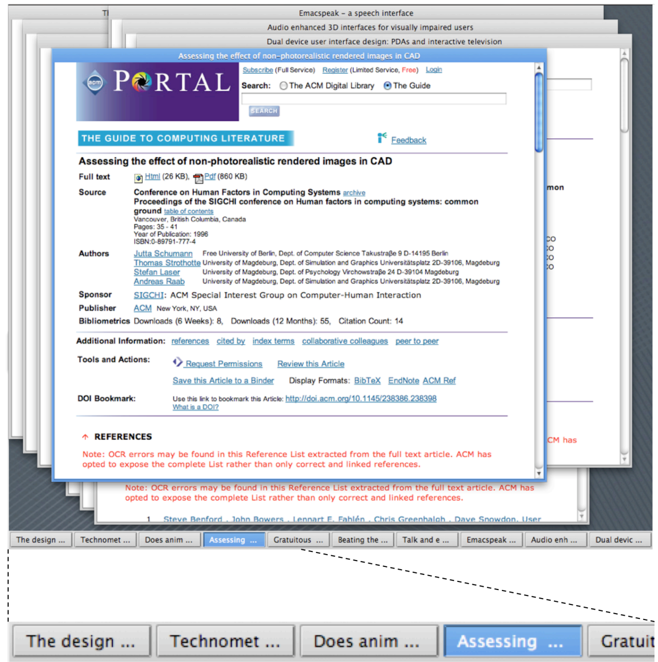

Figure 3: The Overlapping interface. The bar at the bottom of the display contains a tile for each document.

brings the document on top of other documents. Tiles appear from left to right in the bar in the order the documents were opened. All tiles have the same size and all fit in the bar, but only about 10 characters of a document's title is shown. A tooltip with the entire title appears if the mouse cursor hovers over a tile.

Rationale: This interface is included in the experiment because overlapping windows are standard in most systems and graphical user interfaces of many systems include a task bar. Studies have found that the task bar is often used for switching between windows [19,29]. Hutchings et al. found that 78% of the time people had eight or more windows open and consequently, users may experience problems with using the taskbar because only a short part of the windows' titles are visible [19].

The Tabbed interface

The Tabbed interface (shown in Figure 4) contains one window fixed in the center of the display. The window contains multiple documents represented by equally sized tabs under the title bar (see below). Tabs appear from left to right in the order the documents were opened. If there are more documents than tabs, which are minimum 120 pixels wide, a chevron appears to the right of the rightmost tab. Clicking the chevron opens a drop-down list of the titles of all the documents in the order they were added. Documents not visible in the tabs are at the bottom of the list on a gray background. This tab behavior resembles that of the Safari browser, and drop-down lists are common in tabbed document interfaces.

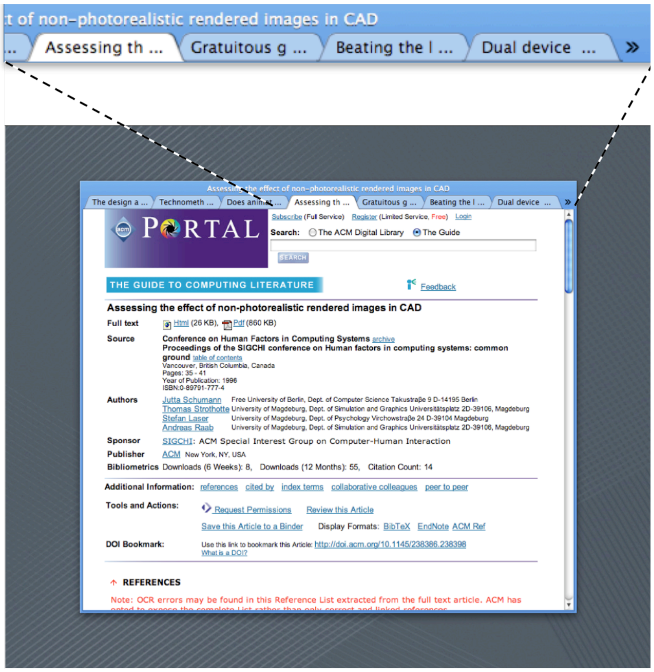

Figure 4: The Tabbed interface. The document tabs below the window title bar contains the tabs for seven of the documents. Clicking the chevron near the right border opens a drop-down list showing all document titles.

Rationale: This interface is included in the experiment because tabbed interfaces are common in widespread web browsers, spreadsheets, and programming environments. Common in many tabbed interfaces is that only a limited number of documents are visible as tabs and users either have to scroll in the tabs or otherwise get a view of the documents that are not shown.

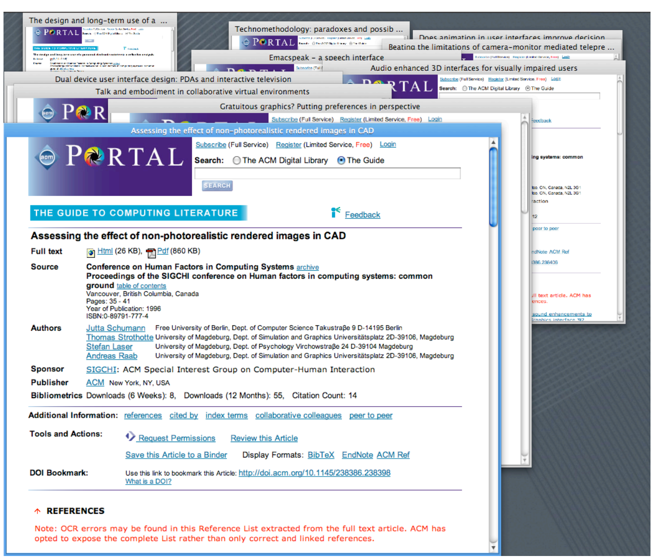

Figure 5: The Piling Recent interface.

The Piling Recent interface

The Piling Recent interface (shown in Figure 5) contains document windows piled in the order that they were last used, with the least recently used document at the bottom of the pile and the currently active document on top. As shown in Figure 6, documents are reduced in size and arranged vertically relative to their position in the pile so that the document at the bottom of the pile is smallest (at 30% size) and nearest the upper display border. Documents are reduced in size by zooming the view of the documents' content. However, the document titles remain in a fixed, readable font size. When documents are opened, their windows are staggered so as to reduce overlap. Also, the staggering may help in remembering documents' location. If the user clicks on a document in the pile, the document moves to the top of the pile and the other documents are rearranged, all with animated transitions.

Rationale: Compared with Overlapping and Tabbed, the Piling Recent interface aims to provide better visibility of the documents, and titles of most documents are readable in full.

The Piling Ordered interface

The Piling Ordered interface (shown in Figure 7) is similar to Piling Recent, except document windows remain in the order they were opened. When clicked, a document is moved from its place to the top of the pile—the document returns to its place in the pile when another document is selected.

Rationale: The spatial layout of document windows in Piling Ordered is more stable compared with Piling Recent, which can potentially help users in revisiting previously used documents.

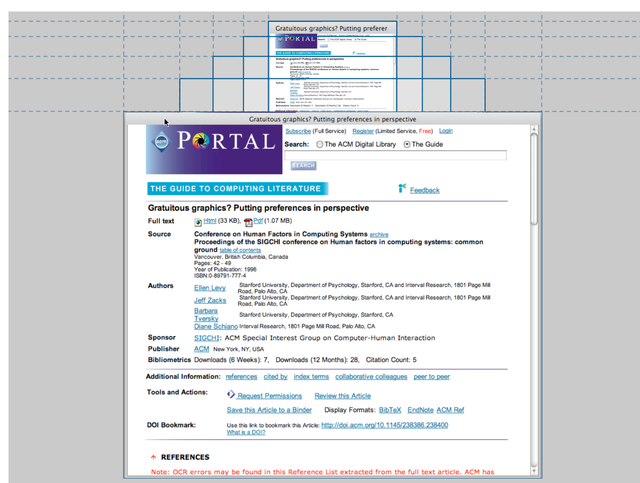

Figure 6: The size and vertical location of a document at each position in a pile of five documents. At the top of the pile, the document is shown at full size.

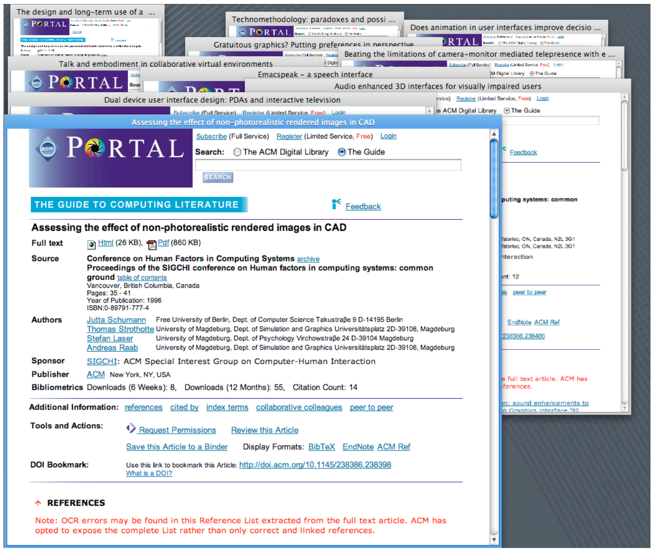

Figure 7: The Piling Ordered interface.

Participants

Eleven participants (eight male) between 25 and 51 years old (M = 34.6) were recruited by word of mouth.

Tasks

Ten web pages were opened in the interface before each task. Pages were selected from among 420 web pages from the ACM Digital Library and showed data on papers published 1994 to 2005. Documents in the experiment were thus all uniform in layout and appearance. Four types of task were used.

Navigate Title tasks required participants to do 25 tasks on finding a particular document, for instance "Click the link to the first author of the document: Pen computing for air traffic control", where the document title varied between eight of the ten documents shown. We wanted to see how spatial layout of pages in the interfaces influenced performance in re-finding documents. Therefore, we used a Zipf-like distribution similar to the one used in Tak et al. [29] for selecting the target documents. The distribution was generated by randomly selecting eight targets. One target was then cued 10 times, one 5 times, then 3, 2, 2, 1, 1, and 1 for the others.

We hypothesized that participants could find the cued document title quickly by scanning in the bar of Overlapping or the tabs in Tabbed.

Navigate Term tasks required participants to find and compare documents with a given term in the title. An example task read: "Click the link to the first author of the document with 'learning' in the title that has the second most recent Year of Publication." Each task of this type used a different set of documents, four of which contained the given term in the title.

We hypothesized that participants could find the term in document titles faster and with less interaction using either Piling interface compared with Overlapping and Tabbed. When document titles were not all shown in full, participants could not determine if the term occurred in a document title directly in the interface and titles were particularly truncated in the tiles in Overlapping and in tabs in Tabbed. In the Tabbed interface, however, participants could call up the list to read the titles in full.

*Navigate All task*s required participants to navigate to every document shown. An example task read: "Click the link to the first author of the document with the highest 'Citation Count'" Each task used a different set of documents.

We hypothesized that participants could traverse the documents quickly in the bar of Overlapping, and also in Tabbed in as far as the documents were visible in tabs.

Compare References tasks asked participants to inspect the 'References' of three of the documents and determine the most referenced author. Each task of this type used a different set of ten documents. The three documents named in the task text had about 15 references each and at least one author referenced in all three documents.

We expected Piling Recent to be suitable for this task because participants could easily find the three documents close together at the top of the pile. We also were curious to see if participants arranged windows for comparison with the Overlapping interface.

These four tasks were chosen to cover basic types of navigation among documents that occur in web browsing or programming. The tasks are similar to tasks used in earlier research [e.g., 20,29]. We are aware that not all user tasks that involve document navigation are among these four types of task. Also, tasks are limited in that they are taken out of context and involve already opened documents.

Materials

Participants used the 1440 x 900 pixel 15" display of a laptop computer. A mouse with scroll wheel was used for input. Tasks were presented in a view in the left side of the display. The interfaces used an area of 1040 x 878 pixels and document windows were 800 x 700 pixels.

Design

The experiment employed a within-subjects design with the factors interface type (Overlapping, Tabbed, Piling Recent, Piling Ordered), and task type (Navigate Title, Navigate Term, Navigate All, Compare References). The order of interface type and task type was systematically varied and counter-balanced across participants. Participants completed 35 tasks with each interface: 25 Navigate Title tasks, four Navigate Term tasks, four Navigate All tasks, and two Compare References tasks.

Procedure

Participants were first given an introduction lasting about 15 minutes, which included instructions on how to use each interface and time for participants to practice tasks similar to the experimental tasks. Participants then completed tasks using each of the four interfaces. After completing the tasks with each interface, participants received a questionnaire about satisfaction with the interface just used. The questionnaire contained nine questions from QUIS [12], two questions from NASA-TLX [17], and two additional questions asking how easy it was to find or re-find documents. Also, participants listed benefits and drawbacks of the interface. After completing all tasks, a final questionnaire asked participants to rank the four interfaces by order of preference. Last, participants were given opportunity to comment on the interfaces. The experiment lasted about an hour and 15 minutes for each participant.

RESULTS

Accuracy

We find no differences in the accuracy with which participants answer tasks, $F(3, 8) = 0.25$, *ns*. On average, participants answered 87% of the tasks correctly, ranging from 88% (with Overlapping) to 86% (with the two Piling variants). We had also expected no difference, as the tasks were relatively easy to solve with all interfaces.

Task completion time

We find an overall interaction between task and interface, $F(9, 2) = 34.93$, $p < .05$, and thus proceed to analyze the time differences on a per task basis (with statistics for significance omitted to save space, all differences are $p < .05$ with linear contrasts or post-hoc tests). Figure 8 shows the task completion times for each interface.

For the Navigate Title task, a post hoc comparison show no difference among interfaces, but to our surprise Tabbed was the fastest interface (M = 5.47s, SD = 2.84), followed by Overlapping (M = 5.66s, SD = 2.30), Piling Recent (M = 5.68s, SD = 2.71), and Piling Ordered (M = 6.21s, SD = 2.54). In particular, the Tabbed interface performed well even for documents that users had to access using the chevron and the drop down list.

For the Navigate Term task, we find that the two Piling interfaces (Ordered: M = 53.3s, Recent: M = 53.1s) are significantly faster than the other interfaces (Overlapping: M = 60.6s; Tabbed: M = 63.8). This confirms our hypothesis that participants could find documents with the term faster using the Piling interfaces because titles of most documents are readable in full. Because the term is not always visible in truncated titles in Overlapping or Tabbed, participants may have spent more time switching between documents or calling up tool tips to look for the term.

For the Navigate All task, we find that Overlapping (M = 21.6s) is significantly faster than Tabbed (M = 35.1s). They are in turn significantly faster than Piling Ordered (M = 39.6s), which is significantly faster than Piling Recent (M = 47.3s). Looking at average task completion times, all participants were faster with Overlapping. This finding agrees with our expectation that participants could traverse the documents by clicking each tile in the bar. One reason participants were slower using Piling Recent compared with Piling Ordered is that the fixed document order in Piling Ordered helped participants revisit documents. Another

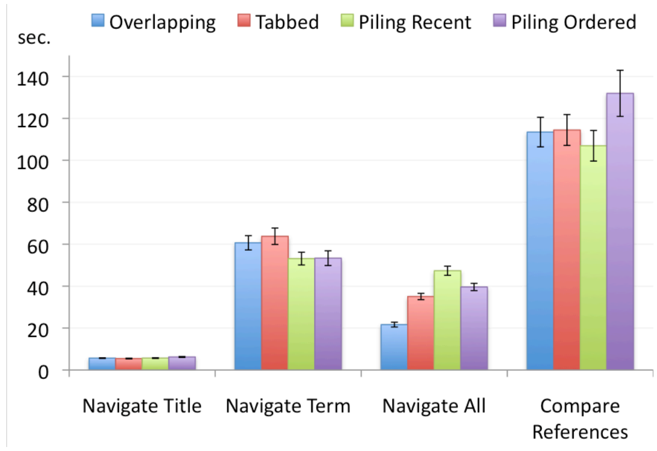

Figure 8: Average task completion times (in seconds). Bars indicate the standard error of the mean.

reason, supported by our informal observations, is that participants sometimes lost track of the first document they viewed in sequentially traversing the documents. Hence, participants may have spent time looking at the same document twice.

For the Compare-references task, we find that Piling Ordered (M = 131.9s) is significantly slower than the other three interfaces (Ms between 113.5s and 116.7s). One reason for this is that participants had trouble remembering the location of the three documents they had to compare references in. Compared to the other tasks involving revisiting documents, this task was cognitively more demanding. It might therefore have helped participants that the three documents could all be found in the order they were last used at the top of the pile in Piling Recent.

Satisfaction and Preference

A multivariate analysis of variance on the 13 satisfaction questions shows no overall differences in satisfaction between interfaces, Wilks' λ= .274, $F(3, 8) = 1.18$, $p > .3$. Only for a question on whether tasks may be completed "with difficulty" or "with ease" do we find a significant difference, $F(3, 40) = 3.20$, $p < .05$, so that the tasks solved with the Overlapping interface are perceived to be easier than tasks with the other interfaces.

In order of preference the interfaces are: Overlapping (M = 1.91, SD = 1.04), Piling Recent (M = 2.27, SD = 1.19), Tabbed (M = 2.82, SD = 1.08), and Piling Ordered (M = 3.00, SD = 1.00). However, an overall analysis of variance on the ranks suggest that these differences are only marginally significant, $F(3, 40) = 2.361$, $p = .086$.

About Overlapping, six participants listed the taskbar as an advantage, three mentioning the arrangement of tiles in a row. Three participants liked that windows could be arranged for comparing their content. Two participants said it was hard to read or get overview of the titles.

About Tabbed, participants described as advantages that it was “clear”, “simple”, and “familiar”. Also, three participants mentioned the alignment of tabs in a row. Six participants expressed uncertainty about how the drop-down list in Tabbed worked, although two participants said it gave good overview of document titles.

For both Piling interfaces, four participants mentioned overview of all documents as an advantage and three participants said the interfaces were good for finding a document. Five participants said about Piling Recent that it was confusing how documents were ordered, but three participants liked having recently used documents on top. Two participants said locations of documents were hard to remember in Piling Ordered, while two participants found it good for revisiting documents. Comments from three participants suggested problems in navigating through all documents using Piling Ordered.

Interaction with the Interfaces

We analyzed data logged during the experiment to help uncover differences in how participants' used the interfaces to complete the tasks. We summed the number of times that participants navigated to a document by clicking either in the document's window, in a tile (only in Overlapping), or in a tab or pop-up menu (only in Tabbed). Table 2 shows the average number of such navigation actions for each interface. Also, we summed the distances that the mouse pointer traveled between mouse button events or tooltip call-ups (shown in Table 3). Distance was calculated as the diagonal between screen coordinates of the mouse pointer. Because the Navigate Title task involved only single-step navigation, we did not analyze the data for that task.

For the Navigate Term task, participants navigated 10.6 times in average using Overlapping and 8.3 times using Tabbed. In contrast, participants navigated 7.2 and 5.7 times in average using Piling Recent and Piling Ordered. This suggests that participants were faster with the Piling interfaces compared with Overlapping and Tabbed because they had to navigate to fewer documents to answer the task. Participants navigated to fewer documents with Tabbed than with Overlapping, which might be because some participants used the menu to find documents with the given term. It is not clear why participants navigated to more documents with Piling Recent compared with Piling Ordered. We wonder if the changing order of the documents made it difficult to revisit the correct document. The mouse pointer travelled relatively less between mouse events with Overlapping compared with the other interfaces, which indicates the benefit of traversing all the document tiles in the bar.

	Navigate Term	Navigate All	Compare References
Overlapping	10.6	11.4	5.7
Tabbed	8.3	12.3	5.0
Piling Recent	7.2	13.9	5.2
Piling Ordered	5.7	11.9	5.9

Table 2: Average number of times that participants navigated to a document.

	Navigate Term	Navigate All	Compare References
Overlapping	75	32	123
Tabbed	109	53	84
Piling Recent	111	192	99
Piling Ordered	105	121	94

Table 3: Average distance that the mouse pointer traveled between mouse events for each interface (100=average distance for task across interfaces).

For the Navigate All task, participants navigated in average to two more documents with Piling Recent than with Piling Ordered. This difference supports the explanation that participants sometimes lost track of the first document they viewed when using Piling Recent, resulting in slower performance. Looking at the mouse pointer travel distance, it seems Overlapping and Tabbed required less user efforts (32 and 53 in average distance) in navigating through all documents than the Piling interfaces (192 and 121 in average distance). This supports the explanation that the aligned tiles and tabs help scanning through the documents.

For the Compare-references task, participants navigated to between 5 and 6 documents in average with all interfaces; which are surprisingly few switches between documents considering that participants had to compare about 15 references in each document. Piling Ordered was slower maybe because it was be difficult to remember the location of the three documents. Participants navigated only slightly more with Piling Ordered than with Piling Recent. It is unclear why the mouse pointer travel distance was larger with Overlapping compared with the other interfaces. It might be because some participants rearranged the documents to better compare them.

DISCUSSION

We emphasize three findings from the experiment, relate them to the literature, and elaborate on task effects.

Visibility. The spatial arrangement in the Piling interfaces makes all documents visible and most titles readable in full. This might explain why participants completed Navigate Term tasks faster using the Piling interfaces. In contrast, document titles are truncated in Overlapping and Tabbed so participants had to inspect each tab or tile, either by using tooltips, by clicking the tab or tile to bring up the document, or by using the drop-down list in Tabbed.

We note that the web pages used in the experiment were uniform in layout and appearance. However, it seems likely that the visibility in the Piling interfaces would improve navigation between visually distinct documents, for instance because users recognize content. Mander et al. studied piles of documents with varied content, but their piles were tightly arranged and used thumbnails [23].

Layout. The aligned layout of tiles and tabs in Overlapping and Tabbed allow for scanning document titles and for clicking through each document in sequence, with little mouse movement. Participants thus completed Navigate All tasks faster using Overlapping and Tabbed compared with the Piling interfaces, which seemed to require more efforts in navigating through all documents. In Piling Recent, participants could continually click the document nearest the upper edge of the display to bring it on top. The staggered placement required moving mouse horizontally, however, and to some participants it seemed unclear when all documents had been viewed. In Piling Ordered, participants had to move the mouse to each document in the pile to bring it on top.

We were surprised that the Tabbed interface performed well, despite the efforts required in navigating to documents not visible in the tabs. Still, only ten documents were used in tasks, and we wonder how tabbed interfaces scale to more documents. Also, participants were effective at navigating many windows in sequence using the taskbar. However, in real life the taskbar often contains different application windows and documents are sometimes collapsed into a single tile.

Further, while techniques for leafing through many documents in tabbed windows [6] and in piles [4] promise to support looking sequentially through documents, we have yet to see empirically studies of their use. Our data suggest that this might be particularly useful for piling interfaces.

Stability. Piling Ordered was slower in Compare References tasks, which suggests that participants had difficulties remembering where the documents were. In contrast, Piling Recent seems more useful in this task because the three recent documents all were near the top of the pile. Tak et al. reach a different conclusion in a study of window switching interfaces: they found that a spatially stable layout allows faster switching to windows often switched to than a recency layout [29]. The importance of stability and recency in the spatial layout clearly depends on the task performed.

Moreover, in this study participants navigated only between documents that were unknown to them. The usefulness of piles in real use situations, where users open and arrange documents themselves, remains to be further researched. In particular, we are curious if spatially stable piles of documents, which were inferior in this experiment, might be useful for structuring work with multiple documents.

CONCLUSION AND FUTURE WORK

Overlapping windows and tabbed document interfaces are widely used in desktop environments, but they provide limited support for navigating among documents. Aimed at providing such support, this paper has investigated the use of automatically arranging documents in piles. Based on five dimensions, we have explored the design space of interfaces for navigation among documents using piles.

In an experiment with 11 participants, we have compared four interfaces that support navigation among documents: one using overlapping windows, one using a tabbed window, and two that arranged documents in a pile. Strong tasks effects were found in task completion times. In tasks that used visual features of documents, participants performed significantly better using piles, likely because visibility of documents was better compared to overlapping and tabbed windows. In contrast, participants performed better using overlapping and tabbed windows in tasks that required participants to scan through all documents. The aligned layout of document titles in the task bar and the tabbed windows' tabs provided for scanning titles in sequence without much effort. Overall, participants preferred overlapping windows, commenting the taskbar as an advantage. In conclusion, results suggest that automatic spatial arrangement of documents in piles provides visibility of documents useful in some tasks, but suggest that techniques for aligning documents could be useful in tasks that involve serial scanning.

For future work, several questions about the use of piles in interfaces for supporting navigation remain to be addressed. First, to provide benefits similar to those enjoyed from arranging and piling documents on desks, more work is required that looks beyond navigation among documents in a single pile. We have begun exploring interaction techniques that allow users to effortlessly arrange documents side by side, in multiple piles (see Figure 1). Another issue concerns how users will arrange and pile documents in focal and peripheral regions, for instance by arranging and ordering documents in piles in the periphery to structure their tasks and piling documents in the focus that are frequently used. This requires empirical studies of people working on more complex tasks than the simple navigation tasks used in the present study.

Also, participants in this study only navigated among 10 documents, with limited screen size. We wonder how piling interfaces scale compared to widespread techniques in navigating among many documents and with large displays.

Last, it remains unclear how transient techniques can use spatial arrangement and representation of documents to support navigation among documents. We will address this in future work and explore how transient techniques can be used for leafing through document piles.

REFERENCES

1. Exposé. Apple Computer, Inc.: http://www.apple.com.
2. Windows Taskbar. Microsoft: http://www.microsoft.com.
3. Reinventing Tabs In the Browser, Mozilla Labs Design Challenge 2009. http://design-challenge.mozillalabs.com/summer09/. Retrieved 2010 May 9.
4. Agarawala, A. and Balakrishnan, R. (2006), Keepin' it real: pushing the desktop metaphor with physics, piles

and the pen, in *Proceedings of CHI '06*. ACM: New York, NY, USA. p. 1283-1292.
5. Aliakseyeu, D., Subramanian, S., Lucero, A., and Gutwin, C. (2006), Interacting with piles of artifacts on digital tables, in *Proceedings of AVI '06*. ACM: New York, NY, USA. p. 159-162.
6. Beaudouin-Lafon, M. (2001), Novel interaction techniques for overlapping windows, in *Proceedings of UIST '01*. ACM: New York, NY, USA. p. 153-154.
7. Bernstein, M. S., Shrager, J., and Winograd, T. (2008), Taskposé: exploring fluid boundaries in an associative window visualization, in *Proceedings of UIST '08*. ACM: New York, NY, USA. p. 231-234.
8. Bi, X. and Balakrishnan, R. (2009), Comparing usage of a large high-resolution display to single or dual desktop displays for daily work, in *Proceedings of CHI '09*. ACM: New York, NY, USA. p. 1005-1014.
9. Bly, S. A. and Rosenberg, J. K. (1986), A comparison of tiled and overlapping windows. SIGCHI Bull., 1986. 17(4): p. 101-106.
10. Bury, K. F. and Darnell, M. J. (1986), Window Management in Interactive Computer Systems. SIGCHI Bull., 1986. 18(2): p. 65-66.
11. Card, S. K., Robertson, G. G., and York, W. (1996), The WebBook and the Web Forager: an information workspace for the World-Wide Web, in *Proceedings of CHI '96*. ACM: New York, NY, USA. p. 111-117.
12. Chin, J. P., Virginia, A., and Norman, K. L. (1988), Development of an instrument measuring user satisfaction of the human-computer interface, in *Proceedings of CHI '88*. ACM Press. p. 213-218.
13. Czerwinski, M., Smith, G., Regan, T., Meyers, B., Robertson, G., and Starkweather, G. (2003), Toward characterizing the productivity benefits of very large displays, in *Proceedings of Interact 2003*. p. 9-16.
14. Czerwinski, M., van Dantzich, M., Robertson, G., and Hoffman, H. (1999), The Contribution of Thumbnail Image, Mouse-over Text and Spatial Location Memory to Web Page Retrieval in 3D, in *Proceedings of Interact '99*: Edinburgh, UK.
15. Dubroy, P. and Balakrishnan, R. (2010), A study of tabbed browsing among mozilla firefox users, in *Proceedings of CHI '10*. ACM: New York, NY, USA. p. 673-682.
16. Grudin, J. (2001), Partitioning digital worlds: focal and peripheral awareness in multiple monitor use, in *Proceedings of CHI '01*. ACM: New York, NY, USA. p. 458-465.
17. Hart, S. G. and Staveland, L. E. (1988), Development of NASA-TLX: Results of empirical and theoretical research, in *Human Mental Workload*, P.A. Hancock and M. Meshkati, Editors. Elsevier: Amsterdam, The Netherlands. p. 139-183.
18. Henderson, J., D. Austin and Card, S. (1986), Rooms: the use of multiple virtual workspaces to reduce space contention in a window-based graphical user interface. ACM Trans. Graph., 1986. 5(3): p. 211-243.
19. Hutchings, D. R., Smith, G., Meyers, B., Czerwinski, M., and Robertson, G. (2004), Display space usage and window management operation comparisons between single monitor and multiple monitor users, in *Proceedings of AVI '04*. ACM Press: New York, NY, USA. p. 32-39.
20. Kandogan, E. and Shneiderman, B. (1997), Elastic Windows: evaluation of multi-window operations, in *Proceedings of CHI '97*. ACM: New York, NY, USA. p. 250-257.
21. Ko, A. J., Aung, H., and Myers, B. A. (2005), Eliciting design requirements for maintenance-oriented IDEs: a detailed study of corrective and perfective maintenance tasks, in *Proceedings of ICSE '05*. p. 126-135.
22. Malone, T. W. (1983), How do people organize their desks?: Implications for the design of office information systems. ACM Trans. Inf. Syst., 1983. 1(1): p. 99-112.
23. Mander, R., Salomon, G., and Wong, Y. Y. (1992), A ``pile'' metaphor for supporting casual organization of information, in *Proceedings of CHI '92*. ACM: New York, NY, USA. p. 627-634.
24. Miah, T. and Alty, J. L. (2000), Vanishing Windows--a technique for adaptive window management. Interacting with Computers, 2000. 12(4): p. 337-355.
25. O'Hara, K. and Sellen, A. (1997), A comparison of reading paper and on-line documents, in *Proceedings of CHI '97*. ACM: New York, NY, USA. p. 335-342.
26. Oliver, N., Czerwinski, M., Smith, G., and Roomp, K. (2008), RelAltTab: assisting users in switching windows, in *Proceedings of IUI '08*. ACM: New York, NY, USA. p. 385-388.
27. Robertson, G., Horvitz, E., Czerwinski, M., Baudisch, P., Hutchings, D. R., Meyers, B., Robbins, D., and Smith, G. (2004), Scalable Fabric: flexible task management, in *Proceedings of AVI '04*. ACM: New York, NY, USA. p. 85-89.
28. Robertson, G., van Dantzich, M., Robbins, D., Czerwinski, M., Hinckley, K., Risden, K., Thiel, D., and Gorokhovsky, V. (2000), The Task Gallery: a 3D window manager, in *Proceedings of CHI '00*. ACM: New York, NY, USA. p. 494-501.
29. Tak, S., Cockburn, A., Humm, K., Ahlström, D., Gutwin, C., and Scarr, J. (2009), Improving Window Switching Interfaces, in *Proceedings of INTERACT '09*. p. 187-200.
30. Voida, S., Mynatt, E. D., and Edwards, W. K. (2008), Re-framing the desktop interface around the activities of knowledge work, in *Proceedings of UIST '08*. ACM: New York, NY, USA. p. 211-220.
31. Weinreich, H., Obendorf, H., Herder, E., and Mayer, M. (2008), Not quite the average: An empirical study of Web use. ACM Trans. Web, 2008. 2(1): p. 1-31.
32. Zhang, X., Qu, Y., Giles, C. L., and Song, P. (2008), CiteSense: supporting sensemaking of research literature, in *Proceedings of CHI '08*. ACM: New York, NY, USA. p. 677-680.