Piles Across Space
Breaking the Real-Estate Barrier on PDAs

Tony Hsieh
Stanford University
Department of Computer Science
Stanford, CA 94305, USA
thsieh@stanford.edu

QianYing Wang
Stanford University
Department of Communication
Stanford, CA 94305, USA
wangqy@stanford.edu

Andreas Paepcke
Stanford University
Department of Computer Science
Stanford, CA 94305, USA
paepcke@cs.stanford.edu

ABSTRACT

We describe an implementation that has users 'flick' notes, images, audio, and video files into piles beyond the screen of the PDA. This scheme allows the PDA user to keep information close at hand without sacrificing valuable screen real estate. It also obviates the need to browse complex file trees during a working session. Multiple workspaces can be maintained in persistent store. Each workspace preserves one configuration of off-screen piles. The system allows multiple PDA owners within ad hoc radio range to share off-screen piles. They point out to each other where a shared pile is to reside in space. Once established, all sharing partners may add to the pile and see its contents. One application is to support biodiversity researchers in the field, where they generate data on their PDA and need to keep it organized until they return to their field station. We conducted an experiment where participants used our system with up to ten simultaneous piles. Not only were they able to operate the application, but they remembered the location of piles when placed in different physical environments and when asked to recall the locations several days after the experiment. We describe gender differences that suggest particular design choices for the system.

ACM Classification H5.2 [Information interfaces and presentation]: User Interfaces. - Graphical user interfaces.

General Terms Design, Experimentation, Human Factors

KEYWORDS: mobile computing, PDA, piles, information management, information sharing, screen real-estate

INTRODUCTION

The usefulness of Personal Digital Assistants (PDAs) is consistently compromised by their meager screen real estate. Many applications require the mobility and light weight of the PDA format, yet suffer from the unavoidably small display size.

We most recently encountered this problem while designing handheld information collection gear for biodiversity researchers. They travel to sometimes isolated places and spend their days walking in the field. On the way they need to collect or generate data. The information might be audio recordings of birds, photographs, or observational notes. This information needs to be managed on the PDA without fuss.

A more subtle generation of information occurs when these researchers work on in-field species identification. Data structures similar to decision trees (dichotomous keys) are used for this purpose. Each node in the tree corresponds to a characteristic that is common to a subset of species, such as the color of their blossoms, or the shape of their stems. The values of these characteristics in the specimen under consideration guide the researcher's travel downwards in the decision tree until a terminal is reached, signalling the successful identification of the specimen.

Unfortunately, it is not always possible to determine all of a given specimen's characteristics precisely. The plant might not currently be in bloom, or the stem is damaged. This uncertainty requires exploratory travel down the tree. Each decision at an uncertain node is a hypothesis that may need to be revisited, either by backtracking the tree traversal while still in the field, or through eventual consultation with an expert at the field station. What is therefore needed is a way of easily managing the accumulating hypotheses.

One approach to this challenge is to fall back on the traditional folder/file metaphor. Researchers would open a file browser, access a file of notes, and append information. In the case of visual or audio media each information item would be stored in a file. Another approach might be bookmarks to which the researcher could return. These might manifest on the screen as tabs, or be accumulated on a separately activated screen. A third solution might be an always available audio recorder where the researcher could note hypotheses or other observations. All such solutions retain a degree of inconvenience. Some, like the file folder approach distract from the focus of the activity at hand. Others, like the bookmarks accumulate very quickly, because they are fine grained organizational schemes. Audio notes are slow to search.

We have been exploring a different route, which was inspired by four existing bodies of work. The first is the oft explored spatial memory of which humans are capable [8, 25]. The second is the notion of 'halos' as published in [2]. The third is the ancient method of constructing *memory palaces* to remember facts and objects, and the fourth is the observation that piles of materials are a frequently found organizing principle [21].

Spatial Memory

Most of the common user interfaces rely on spatial, as well as muscle memory. For example, most applications place their 'File' menu on the upper left of the screen. After extended use, the operators of these applications will not even verify their menu choice by reading its name. Their spatial memory guides them towards the quick selection of the correct choice.

Muscle memory is essential for flow menus [14] that have users move their mouse outward from a starting point on the screen. As they follow a radial motion, the system reveals new options that will have the user move the mouse on a circular path around the starting point, until further radial movement is triggered by a new option. It has been observed that users learn frequently followed paths through these menus, turning their movements into a fluid gesture, rather than read—navigate sequences.

Halos

The notion of 'halos' was introduced as on-screen placeholders for information that is available off-screen. The original application was the browsing of maps on a PDA display. The off-screen information in that case were points of interest, like restaurants, that were located on the map beyond the confines of the screen.

Halos result from drawing an imaginary circle around the off-screen point of interest. The circle is dimensioned such that only a small arc is visible on the screen. The authors of [2] explored how well users could judge the distance and direction of the invisible points of interest by inspecting the shapes of their halo's arcs. The shallower the arc, the further away the point of interest would be located from the edge of the screen. While the authors' study found that distance estimates were not very accurate, the presence of the unobtrusive arcs was an excellent indicator of the existence and direction of the off-screen points of interest.

Memory Palaces

As early as the Roman Empire the technique of memory palaces was taught and refined. This method allows its practitioner to remember long sequences of facts and items. Such a capability was needed, for example, by story tellers that propagated historic events and social norms without an ability to write, or reading skill requirements on the part of audiences.

A mental image of a space, such as a palace, a cathedral, or a city is constructed. The facts to remember are then placed in that imaginary space. Subsequent retrieval works by mentally retracing one's steps through the space.

Piles

Piles, finally, are used by many to organize materials into topically related collections.

Figure 1: Off-screen imaginary piles hold sorted photos.

Piles on PDAs

Our *piles across space* (PAS) implementation allows users of PDAs to 'flick' information, such as photos on their screen, towards the edge of the PDA with their stylus. The system animates the flicked information to follow the trajectory that the flicking motion implies. The information appears to float off the display, seemingly coming to rest in space beyond the edge of the screen. As the information floats beyond the visible portion of the viewing plane, it leaves behind a halo to remind the user of the information's continued existence. As users flick more information items towards that halo they will observe the arc slowly increasing in thickness to indicate that all of the flicked information is accumulating on an off-screen pile. Information other than photos can include audio recordings, text, video clips, decision tree nodes, measurements, or any other locally generated online materials.

Figure 1 shows an example. Two imaginary piles hover beyond the left edge and the upper right corner of the display.

We have implemented three extensions to this notion. The first is to share imaginary piles with other PDA operators nearby. The second is to flick items not to an off-screen pile, but to a large display in the room, and the third is to imagine the flicking to submit information not to a pile or display, but to a processing service.

Pile Sharing

Towards illustrating the sharing of piles, imagine two siblings, Frank and Jane, each sorting a different collection of their family's photos. Both are in the same room, and both are using a separate PDA for their work. They organize the images into off-screen piles, dedicating each pile to images of one family trip. Frank might now point out to Jane the location of a physical piano in the room. They agree on the piano to be the imaginary location of the Paris pictures that occur in their two collections. From then on both of the siblings flick Paris photos that they come across towards the piano. They are now sharing the Paris pile. Whenever either of them examines the pile, they see the union of the Paris photos that have been flicked to the pile so far.

Of course, the location of the piano in the room relative to the two PDA screens is different. For Frank the piano might be positioned beyond the upper right corner of his screen, for Jane the same piano might be located beyond the upper left of her screen. For the system to remain intuitive to both, each must be able to flick Paris photos towards the piano as it appears from their perspective. Upper right for Frank, upper left for Jane. Furthermore, once the two go off to their respective apartments across town, Jane should be able to retrieve the Paris photos from the upper left corner of her screen, while Frank should be confident to find Paris on the upper right.

Piles and Large Displays

The second extension involves a wall mounted or large table top display. PDA users can enter the room and flick individual pieces of information, or entire piles to the communally visible display surface. The material can be discussed or manipulated on the display, and might then be flicked back to one or more PDAs in the room. If this facility were available beyond our prototype, biodiversity researchers, for example, could return to their field station and conveniently consult with the then available group of experts or fellow students.

Piles as Service Queues

The third and final extension we implemented is to make one or more piles be the input queues to services. For example, in a wirelessly connected environment the flicking of text to the lower right corner of the screen might submit the text as a query to Google. Alternatively, terms flung in that direction might trigger a lookup in a locally stored reference work.

Space limitations restrict us here to discussions of the single and multiple user implementations.

Challenges

The very simple organizing facility of piles across space is quick to understand and operate. The flicking motion fits naturally into the workflow, which on PDAs requires the stylus to be near the display surface much of the time. In spite of its simplicity, the full design of the piles across space concept poses a number of challenges. For example, how should the piles be recalled to the screen? How can one enable users to save coherent chunks of their work for later retrieval or resumption of work? How can the contents of piles be browsed quickly? Does the information model allow items to be part of multiple piles at the same time?

The most fundamental question, however, is whether humans are able to operate the system effectively. In particular, we need to find out just how many piles an operator can keep track of, whether the required spatial memory is retained long enough between usage sessions to be of practical use, and how much these two parameters are impacted by age. We conducted an experiment that explored these questions for up to 10 piles and an age range of 18–56.

We begin in the following section with a description of our prototype, and the details of the information and interaction models for both single and multi user deployment. We then continue with a description of our experiment setup, its results, and our discussion of the outcome and its implication for further PAS design work. From there we move to drawing attention to additional related work, and an outline of steps that we plan to take next.

SYSTEM DETAILS

We began with the absolute minimal feature set and then reluctantly added to the set as needed. The following subsections introduce the resulting facilities.

Single User Operation

We call the screen layout of Figure 1 a *workspace*. We call the four halos in the Figure *pile markers*. As will be relevant in the evaluation section, more than four piles may be part of one workspace. The system retains both the composition of items that constitute each pile, and the layout of piles in a workspace. The central area of the screen is called the *work area*. It contains the information that is currently the focus of the user's attention.

Users may accumulate any number of workspaces, all of which they may recall. At any given time, however, only one workspace is active. Recalling a workspace restores the constituent piles to their most recent positions, and recreates the layout of items in the work area.

Populating the Work Area Piles contain *items*, which may be of different media types. In our implementation the item types are photos, video clips, audio files, text files, and PDA-sized sheets with hand-drawn electronic ink. Users populate the workspace with items in several ways: one may open a standard directory of the underlying Windows system. This action arranges media-specific icons in a grid on the screen. The grid is vertically scrollable, if necessary.

A second method for populating a workspace is the production of information on the PDA itself. For example, recordings on PDAs that are equipped with a microphone, photographs that are shot with a built-in camera, or electronic ink notes the user produces with a linked-in application. We launch applications from within a workspace as necessary.

A third way to acquire items is to share piles that were created by other users (see below). A fourth acquisition path, finally, is to have items flicked to the current workspace from our top projection DiamondTouch table surface.

Stylus Interactions Possible interactions with an active workspace are flicking, as described in the Introduction, dragging, single- and double taps with the stylus, and sustained pressing of the stylus on one screen spot.

The action that single taps elicit depend on the item that is tapped. A tap on an audio file starts playback, another tap stops the audio. A tap on a text file raises a tooltip with the first sentence of the file. The choice of the first sentence is a placeholder. A better choice would be a meaningful summarization.

Double-tapping an item activates it as one would expect in a Windows environment. For example, double tapping an audio file opens a player application. Similarly for video and ink notes. Double-tapping on an image enlarges it to full screen size.

The action of pushing the stylus pointer onto the screen and holding it there raises a pop-up menu as is common in PDA-based Windows implementations. The possible actions that are available on the menu depend on the context. In some cases the menu offers the copying of an item or pile. When appropriate we use this action to raise a tooltip instead of a context menu.

Simple dragging allows free form arrangement of items on the screen. More interestingly, dragging from within a halo towards the center of the screen 'pulls' the corresponding pile into the work area. The items that currently occupy the work space are displaced by the items in the pile that is being pulled. The displaced items are not deleted. They are absorbed by a special facility in the data model, which we describe next.

Data Model The new location of the displaced items is either the pile from which they were pulled into the work area, or the special *mother pile* at the bottom of the screen in Figure 1. That special pile holds *free* items, that is items that have not yet been flicked to any pile. It is a kind of organizationally neutral ground.

Most commonly, the mother pile contains all items that were brought in by opening an underlying (Windows) directory. Dragging the mother pile to the active area populates it with those items. As the user flicks items to piles they are removed from the mother pile, though not from the underlying directory.

If free items are in the active work area and an existing pile is dragged into the work area, the free items return to the mother pile. If the replaced items are part of a pile they return there when they are displaced.

If a new directory is opened while free items still remain in the mother pile from an earlier 'open' operation, the result is the union of the new and old items.

Items may belong to multiple piles at the same time. Any item deletion from a pile only affects the item's membership of that pile. We have so far not exposed file system deletion in PAS.

Managing Piles and Workspaces The system currently provides two methods for viewing all of the piles and workspaces that are resident on one's PDA. There is no facility to view the 'universe' of piles that belong to others and are not explicitly shared with the viewer's PDA (see sharing below).

The first view, a *pile strip* is a simple strip of pile icons that can be scrolled across the screen. All piles on the local PDA are represented in the pile strip. The strip is activated through the contact menu in any of the pile markers. The purpose of the pile strip is to select a new pile for the pile marker in which the strip was activated. Tapping on one of the pile icons causes the respective pile to replace the pile that is associated with the marker.

Each pile in the strip can be browsed in place by touching the pile's icon with the stylus and then, without lifting the stylus, dragging around a circle on the screen. The result of this action is that the representation of items in the pile successively appear on top of the pile icon.

In the case of photographs, for example, the rotary motion causes one thumbnail after the other to be displayed. This method is similar to the well-studied technique of rapid serial visual presentation (RSVP) that flashes information at the viewer in rapid succession. Two differences between our facility and basic RSVP are that (*i*) the user's speed of rotation controls the speed of the frame changes, and (*ii*) our users can rotate the stylus in both directions and thereby control the order of frame sequencing.

The second technique that supports the management of piles and workspaces is the *workspace history* view. This view displays a series of the user's previously constructed workspaces. Each display sketches the workspace's pile locations. Here, as in the pile strip view, the user can rotate the stylus on the display and thereby control the successive display of items in a pile. Double tapping selects the chosen workspace, which replaces the one that is currently active. The displaced workspace remains accessible through the workspace history view.

Multiple Users

In addition to organizing personal information spaces into off screen piles, PAS enables piles to be shared. The current implementation uses ad hoc wireless (WiFi) networking for inter PDA communication. Ad hoc is different from the more common 'infrastructure' method of using the WiFi facility around town. The latter causes the connecting source to locate a physically nearby apparatus, called an access point. The access point receives the WiFi signal and bidirectionally patches it through to the Internet. The ad hoc network method of using the same equipment at the source merely locates other WiFi clients nearby and tries to connect to them. This allows multiple devices to connect to each other without an access point. Our measurements with the iPaq built-in WiFi yielded a line-of-sight distance limit of around 300ft (about 90m).

Pile sharing is initiated by the person who has created the pile. The user activates sharing from a context menu. This action causes a popup dialog to open on all PDAs within ad hoc range. Each PDA owner has the option of accepting or rejecting the pile. Everyone who agrees to the sharing arrangement then 'sees' the pile. As part of the acceptance action, the recipient user is free to place the newly shared pile anywhere along the edge of the screen. This is how the piano scenario of the Introduction is realized.

Once a pile is shared with a set of participants, all participants can add and remove items from the shared pile. Each such action is broadcast and executed by all participants.

However, even though every participant can act on the pile, the original pile is considered to be the *master pile*. All additions or removals are atomic transactions, so there is no notion of serializability and concurrency control. Remember that the scenario is of a few people working in reasonably close proximity. We do not envision thousands of PDAs to share piles through this mechanism.

Once a participant physically moves out of radio range from the others, the issue of integrity control is more pressing. We want to allow each participant to continue working on the shared pile while out of range. This permissiveness obviously leads to inconsistent copies. Our design decision for this problem is an example of the design minimalism that guided us. Rather than introducing a complex reconciliation protocol, our strategy uses optimistic concurrency control with a centralized master copy. The following occurs when the distant participant reenters radio range with the pile owner.

Both parties see a dialog box on the screen, informing them of the conflict. Each has the choice of accepting or rejecting the other's pile version.

If both decline, the pile sharing arrangement is terminated. Each party retains its own version. If the owner of the master pile accepts the changes made by the non-owner, the master pile is modified by merging the non-owner's action log with the owner's log. If the non-owner then does not accept the owner's changes, the pile sharing does not occur. The owner has ultimate control. If the non-owner wants to participate in pile sharing, he needs to accept whatever changes the owner has made. If the non-owner accepts the changes made by the owner, but the owner declines the changes made by the non-owner, the non-owner's PDA will simply copy the owner's current master pile state. The sharing arrangement continues.

Finally, if both parties agree, the non-owner's log is transmitted to the master. The master time-merges that log with its own and brings the two versions to a consistent state. This state then becomes the new master state and is broadcast to the non-owner.

EXPERIMENT SETUP

We recruited 32 paid participants, 14 women, 18 men. Nine of the participants were 18 or 19 years old. Fifteen were 20-36, while 8 were in the 43-56 range. We used a mixed factorial design for the experiment. The physical setup and procedures of our design were as follows.

We installed 10 physical piles of objects in office A, the 'large' room, and 4 piles in office B, the 'small room.' Both rooms were rectangular and empty, save for tables along three walls, and an office chair.

The piles of physical objects were placed on the tables. For example, in the large room we arranged four piles along each of the long sides of the room, and one pile each along the short walls, for a total of 10 piles. Figure 2 shows a schematic view of the large room.

Along the left wall we heaped sporting goods, office supplies, tools, and books. Each pile contained around three to four items. An easel with a painting and brushes stood along one of the short walls and represented art. Along the right-side walls were toys, coffee making equipment, cosmetics, and a refrigerator. On the door was a hook with coats. The chair was placed near the door, as suggested by the top-down camera image of a person in Figure 2. In the small room we placed only 4 piles, one in each corner of the room.

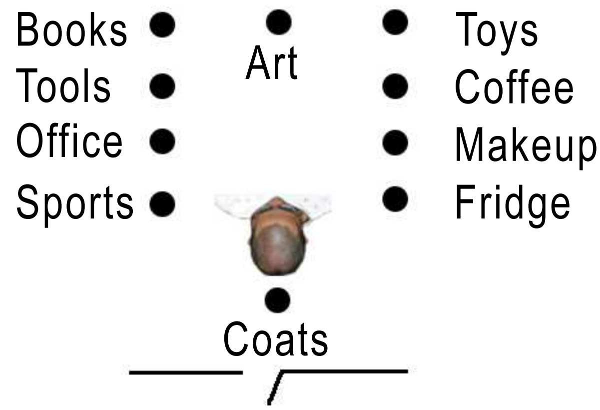

Figure 2: Participant with 10 physical piles on surrounding tables and the coat hanger behind.

Other than the number of piles, the procedure in the large and small rooms were almost identical. We therefore explain below in terms of the large room, mentioning the small room only where needed.

We asked each participant to sit down in the chair, facing a prescribed direction. For example, in Figure 2 the direction is towards the art pile, with the person's back to the door. This initial direction was varied among the participants to counter balance order effects.

Each participant received an iPaq 3500 PDA with ten half circles along the edges (see Figure 3). Each half circle represented one of the physical piles in the room. For each participant the positions of the piles and the positions of half circles on the screen matched at the outset of the experiment.

For example, Figure 3 shows the screen layout that corresponds to the position of the person in Figure 2. The upper left half circle corresponds to the books pile, the half circle on the lower right of the screen corresponds to the refridgerator and the pile of perishable food the appliance implies.

Once the experiment began we successively displayed photos of individual objects on the screen. In the example of Figure 3 this object is an electric drill. The participant was asked to drag the object to the pile to which the item belongs. For the drill in Figure 3 the proper pile is tools, which we see in Figure 2 is one of the piles on the left. The arrow in Figure 3 is not part of the actual PDA screen, but illustrates the movement of the stylus for this electric drill example.

The matches of objects to appropriate piles was simple. We did not introduce any ambiguities. But we asked participants to work as quickly as they could. Successive photos randomly matched different piles. However, each participant ex-

Figure 3: Participants' task is to assign successively appearing images to the correct piles by flicking them off the screen in the proper direction. The Figure shows 2 of the 10 physical piles that surround the participant.

perienced the same sequence of photographs. For the large room condition, this sequence contained 70 photos and for the small room, it had 28 photos.

After the 70 photos were assigned, we asked the participants to play an identical picture matching game on a different PDA. The task was to match two small tiles with same pictures. The game took about 1.5 minutes to complete. The intent of this interlude was to disrupt the short-term memory before the next phase of the experiment.

For the next phase we asked participants to rotate their chair, and thereby the PDA by 90 deg. The participant then matched another 70 photos, after which they played a puzzle again, the chair was rotated by another 90 deg and the participant matched 70 new photos. The last rotation brought the subject to face 270 degs from the original seating position. We did not repeat the initial position.

All participants ran through this basic procedure. However, half of the participants were in what we call the 'fixed' group, while the other half was in the 'relative' group. The different treatment of the two groups modified the correspondence of half circles to physical piles in the room. After each chair rotation we instructed the fixed group to continue using the same half-circle-to-pile mapping that was in force during the initial chair position. The fixed group therefore enjoyed a natural, visual match between half circles and physical piles only for the initial chair position. However, subjects could at any time hold the stylus over any of the half circles to elicit a tool tip that revealed that half circle's association.

In contrast, the relative group's instructions had them reassign the half circles on the screen to the new relative locations of the piles after each turn. The half circles on the screen were, of course, adjusted with each chair rotation to make this new mapping possible. For example, once the participant in Figure 3 rotates by -90 deg, both the top and the bottom of the PDA screen would show four half circles. The broad sides would show one half circle each.

Once participants were finished in the first room, we asked them to accompany us to the street outside the building. They were free to face in any direction, but we asked them to imagine the room they had just occupied, and to tell us where in their mental map the door would be relative to where they were standing outside. Once we noted that position, participants associated another random sequence of photos, this time with no physical clues at all.

Finally, we repeated the experiment by re-entering the building and visiting the room they had not worked in the first time around. If they started with the large room they would continue in the small room after the trip outside, and vice versa. Once the image associations were complete in the second room, participants again stepped outside and worked without physical reference.

After each segment of the experiment we asked participants to fill out a questionnaire. The entire procedure lasted about 1 hour. We balanced both the first exposure to small vs. large room and the first chair position in which participants began within each room. Both of these conditions were in turn balanced for gender.

Our rationale for using 10 and 4 piles, respectively, in the two rooms was initially based on the 7+/-2 rule of thumb for the number of items human beings can on average hold in short term memory. This theory would call for 5 and 9 piles, respectively. However, initial pilot experiments showed that 9 piles did not stretch our test participants' capabilities enough. We therefore expanded to 10 at the high end. The use of five piles for the small room would have introduced a potentially confounding factor: asymmetry. We feared that non-symmetric arrangements would favor the one-out pile, disrupting our ability to compare participant capabilities strictly by number of piles. We therefore decided on four piles for the small room, arranging the corresponding half circles into the corners of the screen.

Throughout the experiment we measured the amount of time participants required to complete the photo associations in each position, the partial completion times for the photos 1-20, 21-40, and 41-70 (analogously for the small room), the number of association errors, and the number of times participants needed to refer to the half circle association reminder tool tips.

We conducted a follow up four days after each participant completed the experiment. We contacted the participants and asked them for their memory about the piles and how they were located for both the *complex* and *simple* conditions. We noted the accuracy of their replies.

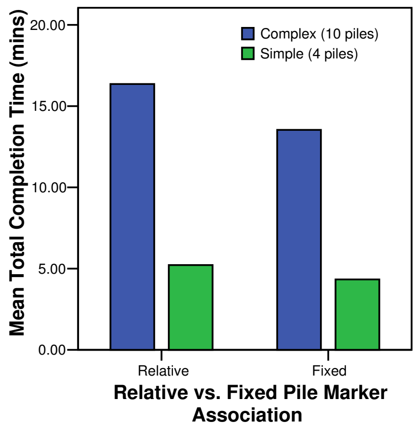

Figure 4: Completion speed for the *relative* and *fixed* pile marker associations. The *relative* setup slows participants down by 17%.

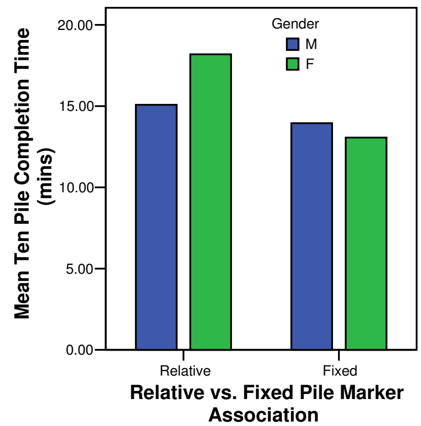

Figure 5: Comparison of how marker association schemes impact the speed of women. Women were 28% faster in *fixed* than in *relative*, while men were not significantly affected.

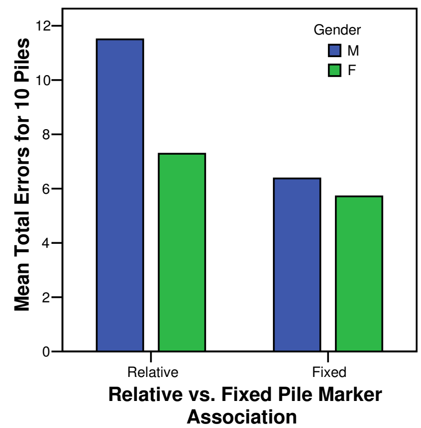

Figure 6: Error rates by gender for the *relative* and *fixed* pile marker associations. The *relative* setup causes a significant increase in errors for males, while the error rates of women were not affected.

RESULTS

The main thrust of our analysis explored how well the participants who operated under the *relative* implementation fared when compared to those who were given the PDAs with the *fixed* correspondence between pile markers and the physical piles in the room. We measured both the speed and the accuracy with which the participants placed each image onto its respective pile.

Within both of these groups we looked for differences in speed and accuracy as the participants dealt with only four piles, the *simple* condition, as opposed to the *complex* condition that had them operate the PDAs with ten piles. We also monitored how often they resorted to the tool tips to remind themselves of which pile marker corresponded to which physical pile.

We then explored whether there were differences in these measurements by age and gender. Participants' long-term memory (LTM) about the piles and the order of piles were recorded and compared across conditions.

Relative vs. *Fixed* Pile Association

Figure 4 shows a comparison of completion times for the *relative* and *fixed* marker correspondences.

Results for both the ten pile *complex* and four pile *simple* setups are shown. The *fixed* marker correspondence enables a 17% faster completion time than the *relative* scheme for both complexities ($t(30)=2.5$; $p<.02$). The *fixed* marker correspondence also led to 38% fewer image assignment errors than *relative* for the ten pile setup. Corrected for unequal variances the effect manifested at $t(25.14)=2.9$; $p<.007$. For just four piles the error rate was not significantly affected by the marker correspondences.

Gender Differences Figure 5 shows that male and female participants were very differently affected by the *relative* method of pile association.

Women were 28% faster in *fixed* than in *relative* with $t(12)=2.4$; $p<.03$. The error rate of women, however, was not significantly affected by the pile association scheme. In contrast, men performed about equally fast with 10 piles in *fixed* and *relative*. On the other hand, male error rates with 10 piles was 45% higher in *relative* than in *fixed* (Figure 6. $t(16)=2.9$; $p<.01$).

Age Differences We compared the nine subjects who were 18 or 19 years old with the eight subjects who were between 43 and 56. Figure 7 shows the mean completion time of the two age groups for the 10 pile setup.

Mature participants were 26% slower than teens ($t(15)=-2.9$; $p<.01$). With four piles the speed difference dropped to 17%.

Long-Term Memory Of our 32 participants, 28 responded to our follow-up request for a sorted list of piles. Combining all errors that these 28 made in the follow-up (i.e. adding errors for their ten pile and four pile setups), we found that those who worked under *relative* had an 81% better recall with $t(26)=2.9$; $p<.03$. In absolute numbers, *relative* par-

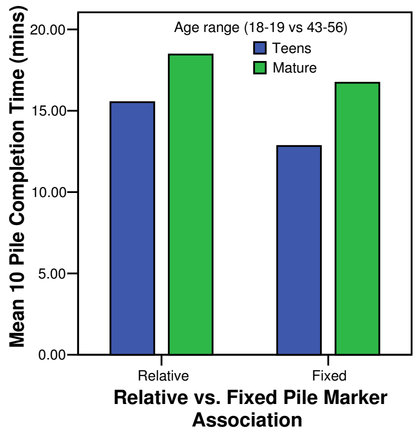

Figure 7: Young participants are faster than mature subjects.

ticipants made .38 errors, while the *fixed* subjects were on average unable to identify and sequence two of the 14 piles. No gender difference was found in long-term memory.

Pile Theme Differences and Tooltips While participants almost never used tool tips, we did observe differences in the average time subjects took to make the association with the various piles. The office pile (a typewriter, stapler, staple remover, scotch tape, and a pair of scissors) slowed the participants down the most. They spent more than 13 seconds on the 35 office-related photos (seven each in the five trials). This compares, for example, with 8sec for coats, 10secs for toys and food. The second most time consuming were tools and books at 11sec.

INFORMAL OBSERVATIONS AND DISCUSSION

The most surprising to us was the ability of all users to work with multiple off-screen piles and rarely be confused. We had originally thought that ten piles would stress participants to their limits. But only 30 errors occurred for the 2,240 images that the 32 subjects assigned outdoors where the physical piles were unavailable and participants needed to rely on their mental model of the pile layout. Of most concern to us are the outdoor results, because they most directly affect our future design choices for PAS. Only five participants ever used the tooltips outdoors to remind themselves of the pile layout, but even they took advantage of this feature less than four times for their respective 70 outdoor images. As to the long-term memory (LTM) test we conducted four days after the experiment, only 17 LTM lapses occurred. Fourteen participants had no LTM errors at all.

On the one hand, the relative pile marker scheme was the worse choice for all participants, but for different reasons. Changing the marker associations to be in consonance with the physical piles slowed women down. Men suffered a steep increase in error rate. But on the other hand, long-term retention benefited from the relative scheme. We speculate that the increased mental investment that is needed to cope with the *relative* scheme is the likely reason for this stronger retention. We point to some applicable cognitive psychology researches in the Related Work section. Another noteworthy observation with regards to participants' LTM is that they could manage working in two different rooms with non matching pile layouts. The results show that our participants were able to separate two different settings well. This finding will help us develop capabilities that allow users to switch pile layouts when they need to switch context in their work activity.

Some results might be useful for designing pile layouts for PDAs. For example, several participants pointed out that the pile *locations* along the sides of the display were more difficult than the ones in the corners of the screen. The door (home of the clothes pile) played an important role for several participants. At least two of them, when outdoors, chose a tree for a *landmark* to represent the door of the room with the physical piles. Others mentally placed the door where it had been relative to them in their most recent position in the room. Others still reported that they picked one of the piles as an origin, placed it mentally when outdoors, and then derived all other positions.

Several design implications arise from these results. We would like to minimize mental effort for the pile users. The results clearly point towards not moving piles around even if the user still physically resides and reorients in the environment that induced the piles. Recall that our target outdoor application audience are biodiversity researchers in the field. The superiority of *fixed* obviates the need for orientation sensing on the PDAs in the context of piling; an advantage favoring unit cost and battery life. On the other hand, we would like to leverage the LTM benefits of the relative scheme. We might be able to derive a solution to this quandary by noting that the informal reports about participants' mental models indicate that they all use one reference point or landmark to reconstruct the entire pile layout. This strategy implies that even participants in *fixed* introduced a notion of relativity once the physical references are removed.

Three design possibilities come to mind for generating the mental effort of relativity without paying the associated cost in time and errors. We could provide users with a method of communicating to the PDA the layout of real rooms that are part of their lives. For example, they might as part of the customization process with their PDA sketch the layouts of furniture in their living rooms and bedrooms. Analogous to the memory palace approach they would inform the PDA of their room choice at the outset of a session. They would then place all their new piles such that the are aligned with the realities of that room, even when they are physically else where. The hope would be that the mental effort of imagining the room while out in the field would help long-term retention. But once this mental effort has been expended, operation would proceed as in the *fixed* condition.

A second possibility is maybe less promising, but still worth exploration in future work. The PDA would ask users to choose a landmark in their physical environment when they begin work. Such a landmark might be a rock on their right. They would then place a symbol of that rock on their screen in alignment with the landmark. As they reorient in space they would operate as in *fixed*, but the symbol would help them rotate the environment in their mind. It is unclear whether this subtlety will suffice to invoke the LTM benefit of *relative*.

A third possibility for achieving higher long-term retention is to let users place their piles and the associated pile markers anywhere along the screen's edge. This conscious effort of deciding positions might be enough to commit their layout to memory. This approach has a second advantage. As reported above, several participants felt that piles on the right and left of the screen were more difficult to manage then piles at the top. We speculate that this difficulty stems from the double occupancy of the prominent center position along the sides of the PDA. At the top of the screen there is only one pile that occupies the center spot between the corners. Along the sides there are two piles vying for the center. These competitors are consequently easy to mix up. Allowing users to place their pile markers manually will enable them to create asymmetry along the PDA's long edges. For example, one of the pile markers could reside close to one of the corner piles. This choice of placement can, of course, be related to semantic similarity among the two piles.

RELATED WORK

Several projects have studied visualization techniques for small screen devices. One technique involves the use of fisheye views [13] and [4]. However these views often result in distortion, which makes it difficult to perform certain tasks. Baudisch & Rosenholtz [2] studied off screen halos with respect to map-based route planning tasks for small screen devices. Other researchers have used techniques involving arrows pointing into off screen space [18] and [7]. This often involves overlaying the arrows onto other content. Another method for manipulating information on a small screen involves the use of special zoom UIs [3] and [20]. None of these techniques have been specifically used for sorting/organizing information on a small screen device.

In addition to different visualization techniques, reseachers have investigated alternative physical interactions to deal with the small screen of handheld devices. These range from contact and pressure to tilt and motion. Rekimoto [23] uses tilting for menu selection and map browsing. Harrison et al. [17], Small & Ishii [27], and Bartlett [1] use tilt sensors to scroll through and select information on a handheld device. In the ComTouch project, Chang et al. [5] use a pressure sensor to translate hand pressure into vibration intensity between users in real-time.

Other solutions include the use of spatially aware displays. Spatially aware displays use a position-tracked display to provide a window on a larger virtual workspace [11] and [12]. In the Peephole project, Yee [29] combines two-handed pen input techniques with spatially aware handheld computers to navigate and manipulate objects in 3-D. These solutions often involve special hardware that is not preinstalled in handheld devices sold today.

Spatial memory is based on the research result that we remember where things are in a location. Spatial cognition is a frequently explored capability [8] [25]. Cockburn and McKenzie [6] studied the effectiveness of spatial memory in 2D and 3D virtual worlds. In the DataMountain project, Czerwinski et al. [9] [24] showed that spatial arrangement of Web pages allows faster and more accurate retrieval both immediately and 6 months later.

Related work with respect to context dependent memory is also available. Context-dependent cues have strong impact on how people encode and retrieve information. Previous research [15] demonstrated the importance of environmental cues for retrieval. Information in long term memory is more accessible if context-dependent cues are present. In our case spatial memory allows users to create a positional mapping between the physical piles in the real world and the virtual piles on the screen.

Others have studied the gender differences in spatial and navigation tasks, discovering that males tend to perform better in these tasks [16] [19] [26]. Devlin and Bernstein [10] showed that males preferred the use of spatial cues and made significantly less errors than females in a computer simulated campus tour. Tan et al. [28] show that these gender-specific navigation benefits come from the presence of optical flow cues and test this on large displays.

CONCLUSION AND FUTURE WORK

We described our prototype of *Piles Across Space*. This facility breaks through the screen real-estate barrier that PDAs impose on their applications. Users flick all kinds of digital information towards the edge of the screen where it disappears as if joining an imaginary pile that is located off-screen. The system works for individual users as well as for small work groups with multiple radio-linked PDAs.

We briefly described our implementation for variations of this basic idea. The first variation was the ability to fling information to a large screen display. The second was the replacement of piles with service input queues. Anything flung to a queue activates the associated service.

We then described a user experiment that explored several design details for the PAS system. The main topic was whether pile markers along the screen edge should reorient as the user faces in different directions. We found that the maintenance of relative positioning slowed participants down and increased their error rate. On the other hand, the relative positioning also improved longterm memory of the pile layouts.

We were surprised that participants were capable of managing 10 piles with relative ease, and that longterm recall even after—in some cases—14 days was quite satisfactory.

Several design implications of our findings point towards future work. One topic are design changes that will combine the speed and accuracy advantages of the fixed pile position approach with the long-term benefits of the more difficult to manage relative approach.

Another topic will be to allow users to construct their own universes of piles. Beyond this, given the success with 10 piles, we are tempted to add a third dimension, namely stacking two layers of piles along the Z-axis.

A third thrust will be to add efficient pile summarization functions. We want users to recall the contents of offscreen piles within a very short period of time.

Overall we are encouraged by our implementation, the results of the experiment, and most of all by the enthusiasm that we observed in our experiment participants and others when we demonstrate the prototype.

REFERENCES

1. Bartlett, J.F. Rock n' Scroll is Here to Stay. *IEEE Computer Graphics and Applications*. 20(3):40-50, May/June 2000.
2. Baudisch, P., and Rosenholtz, R. Halo: a Technique for Visualizing Off-Screen Location. CHI'03 418-488.
3. Bederson, B., Hollan, J. Pad++: a zooming graphical interface for exploring alternate interface physics, Proceedings of the 7th annual ACM symposium on User interface software and technology, p.17-26, November 02-04, 1994, Marina del Rey, California, United States
4. Carpendale, M.S.T., Montagnese, C.A. Framework for Unifying Presentation Space. In Proc. UIST'01, pp 61-70.
5. Chang, A., OModhrain, S., Jacob, R., and Ishii, H. ComTouch: Design of a Vibrotactile Communication Device. in *DIS2002*. 2002: ACM p. 312-320.
6. Cockburn, A., McKenzie, B. Evaluating the effectiveness of spatial memory in 2D and 3D physical and virtual environments in Proc. CHI '03, pp. 203-210.
7. Cox, D., Chugh, J., Gutwin, C., and Greenberg, S. (1998) The Usability of Transparent Overview Layers. In Proc. CHI '98, pp.301-302.
8. Curiel, J.M. & Radvansky, G.A. (1998). Mental organization of maps. In Journal of Experimental Psychology: Learning, Memory and Cognition, 24, 202- 214.
9. Czerwinski, M., van Dantzich, M, Robertson, G., and Hoffman, H. The Contribution of Thumbnail Image, Mouse-over Test and Spatial Location Memory to Web Page Retrieval in 3D. Proc. INTERACT'1999. 163–170.
10. Devlin, A.S., Bernstein, J. (1995). Interactive wayfinding: Use of cues by men and women. Journal of Environmental Psychology, 15, 23–38.
11. George W. Fitzmaurice, Situated information spaces and spatially aware palmtop computers, Communications of the ACM, v.36 n.7, p.39-49, July 1993
12. George W. Fitzmaurice , Shumin Zhai , Mark H. Chignell, Virtual reality for palmtop computers, ACM Transactions on Information Systems (TOIS), v.11 n.3, p.197-218, July 1993
13. Furnas, G. Generalized fisheye views, In Proc. CHI '86, pp. 16-23.
14. Guimbretire, F., Stone, M., Winograd, T. Fluid interaction with high-resolution wall-size displays, Proceedings of the 14th annual ACM symposium on User interface software and technology, November 11-14, 2001, Orlando, Florida
15. Godden, D.R., & Baddeley, A.D. (1975). Context-dependent memory in two natural environments: On land and underwater. British Journal of Psychology, 66: 325-331.
16. Halpern, D.F. (2000). Sex Differences in Cognitive Abilities, 3rd Edition. Lawrence Erlbaum Associates, Inc: New Jersey.
17. Harrison, B., Fishkin, K., Gujar, A., Mochon, C., and Want, R., Squeeze Me, Hold Me, Tilt Me! An Exploration of Manipulative User Interfaces. CHI'98, 17-24.
18. Harrison, B.L., Kurtenbach, G., Vicente, K.J. An Experimental Evaluation of Transparent User Interface Tools and Information Content Evaluation. In Proc. UIST '95, pp.81-90.
19. Kimura, D. (1999). Sex and Cognition. MIT Press: Cambridge, MA. 1–66.
20. Wendy E. Mackay , Guillaume Pothier , Catherine Letondal , Kaare Begh , Hans Erik Srensen, The missing link: augmenting biology laboratory notebooks, Proceedings of the 15th annual ACM symposium on User interface software and technology, October 27-30, 2002, Paris, France
21. Malone, T. How do people organise their desks? Implications for the design of office information systems In Proc. TOIS '83, pp. 99-112.
22. Rekimoto, J. and Saitoh, M. Augmented Surfaces: A Spatially Continuous Work Space for Hybrid Computing Environments In Proc. ACM CHI '99, Pittsburgh, PA, May 15–20 1999. ACM Press. 378-385.
23. Rekimoto, J. Tilting Operations for Small Screen Interfaces UIST96, 167-168.
24. Robertson, G., Czerwinski, M., Larson, K., Robbins, D., Thiel, D., Dantzich, M. Data mountain: using spatial memory for document management, Proceedings of the 11th annual ACM symposium on User interface software and technology, p.153-162
25. Roskos-Ewoldsen, B., McNamara, T.P., Shelton, A.L. & Carr, W. (1998). Mental Representations of Large and Small Spatial Layouts are Orientation Dependent. In Journal of Experimental Psychology: Learning, Memory and Cognition, 24, 215-226.
26. Sholl, J.M., Bartels, G.P. (2002). The role of self to object updating in orientation free performance on spatial-memory tasks. Journal of Experimental Psychology: Learning, Memory, and Cognition, 28(3), 422–436.
27. Small, D., and Ishii, H. Design of Spatially Aware Graspable Displays CHI97 Companion 367-368.
28. Tan, D., Czerwinski, M., Robertson, G. Women go with the (optical) flow In Proc. CHI '03 pp. 209-215
29. Yess, Ka-Ping Peephole Displays: Pen Interaction on Spatially Aware Handheld Computers In Proc. CHI'03, pp 1-8.