

# Multi-User Piles Across Space

*QianYing Wang*  
Stanford University  
wangqy@stanford.edu

*Tony Hsieh*  
Stanford University  
thsieh@stanford.edu

*Meredith Ringel Morris*   *Andreas Paepcke*  
Stanford University   Stanford University  
merrie@cs.stanford.edu paepcke@cs.stanford.edu

## ABSTRACT

We introduce Multi-User Piles Across Space, a technique that allows co-located individuals with PDAs to share and organize information items (e.g., photos, text, sound clips, etc.) by placing these items in shared, imaginary off-screen piles. This technique relies on human capacities to remember spatial layouts, and allows small co-located groups with limited screen real estate to collaboratively manage information. Each participant can use their PDA's stylus to flick information to shared off-screen piles and view their contents. Connections are implemented through ad hoc WiFi. Optimistic concurrency control provides long term data consistency. We also describe an extension that allows PDA owners to transfer information items and piles to and from a tabletop display.

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

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

## INTRODUCTION

Limited screen real estate is both boon and bane of personal digital assistants (PDAs) and cellular telephones. Small size is the key to their convenience, but also an annoyance when a task requires juggling related information that does not fit on the screen all at once.

The most common solution to this quandary is panning. This technique conceptually places all information on a plane, treating the PDA display as a window onto the information. The content of the window is changed by grabbing the information plane and dragging it about. One variation is to make the device location aware, and to physically move the device, keeping the underlying plane constant [14, 6]. This approach, of course, requires hardware to provide the location awareness. Another approach is to very rapidly flash information to the screen (for example, [5]). This technique, called Rapid Serial Visual Presentation (RSVP) has been shown to be effective for text, but it is not relaxing to use. A third approach is to automatically summarize information shown on a small display [3], or to simply require that information is formatted for handheld devices [1].

All these approaches assume that there is a pre-existing body



Figure 1: View of PAS screen with imaginary piles

of information that is to be browsed on the small display. We address here the situation where small device owners generate information during the course of a task, and are trying to organize that information as it comes to life. For example, consider field biologists who would like to carry a small device on which they can take photographs of species they sight, or where they can record measurements of instruments they have scattered on plots of land. Analogously, take a consumer scenario where a camera owner would like to sort photos on-the-fly.

We implemented Piles Across Space (PAS), which allows users to create virtual piles that conceptually hover around the display. These piles can be created dynamically by flicking information items to them with a stylus. We devoted a separate publication to two user studies that examine single-user interaction with the PAS system [13]. These two studies also explored the mental models that participants developed spontaneously to help them remember where in space they 'parked' their piles. Here, we describe an extension of PAS that supports co-located collaboration. We first introduce the basic user interaction with the PAS system. We then describe how multiple PDA holders can share virtual piles. Co-located groups can also move piles and information items to and from a desk-sized tabletop display to enable a back-and-forth workflow between the private PDA space and this public display area. We close with pointers to additional literature.

## BASIC USER INTERACTION

Figure 1 shows the main view of the system. The central area is called the *work area*. In this Figure the icons represent photographs, sound recordings, and a video clip. We also support hand-written ink notes and text files.

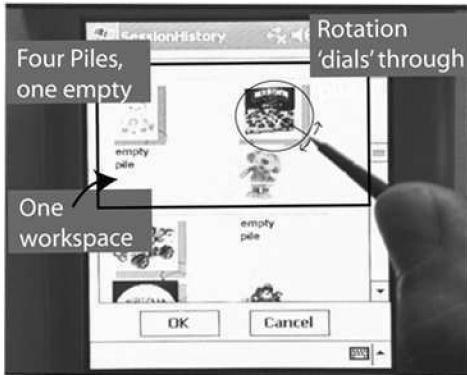


Figure 2: The workspace history browser

Items enter the work area by various means. The items may be created on the PDA itself. For example, on our camera-equipped PDA we can take photos that then display as thumbnails in the work area, which is scrollable. We can also launch processes from within PAS; this facility allows us, for example, to run the ink note application and have the ink notes display as items in the work area.

Another method for placing items in the work area is through imports from the underlying file system. Imports are initiated via a context menu. All imported items are added to a special pile at the bottom of the screen. This pile is called the *mother pile*. In Figure 1 the location is indicated by an up arrow at the bottom of the screen. The arrow indicates that the work area is currently populated with the contents of the mother pile, as opposed to the items from one of the user-made piles.

Users can create piles through a context menu. A new pile manifests as a circle on the screen, which the user grabs with the stylus and places along the edge of the screen such that only a partial arc extends into the display area. Once in place, users can drag items to the new pile marker. In Figure 1 we see four piles, one in each corner. The arcs are called *pile markers*. They remind the user of each pile's location off screen. The system thickens the arc as items are added to a pile. As an additional clue, the number next to the arc indicates how many items are in the pile. We avoid more descriptive labels to preserve a maximum of screen space for the ongoing task.

A pile is pulled into the work area by dragging its pile marker towards the center of the display. Items that are present in the work area at that time are retracted into the mother pile. The work area is then occupied by the pile's items. Items may be copied by dragging them onto the Copy button at the top of the work area. Deletion of items via the context menu removes them from their pile, but not from the underlying folder system. Piles are references to files in the operating system. Items can therefore be part of multiple piles.

As PAS users work on different tasks, they can build multiple pile environments, or *workspaces*. A workspace consists of a populated work area, and a set of piles, along with the locations of these piles. Figure 2 shows an annotated screenshot of this facility.

Each workspace is a context to which the user might wish to return. PAS remembers all the workspaces that a user constructs.

In Figure 2, we see the representation of one workspace at the top of the screen, and a partial workspace at the bottom. The workspace at the top consists of four piles; three contain items, and one is empty. Double-tapping in the white area of the workspace restores that workspace as the active work area. That is, all pile markers will be where the user has left them, and all pile compositions will be restored.

If the images at the top of each pile are insufficient to help the user select the proper workspace, a pile inspection facility can assist. This facility is activated by the user tap-dragging the stylus around a pile as suggested by the superimposed graphic. This motion will continuously shuffle through the pile, exposing all information items in turn.

A shortcoming of our workspace browser is that it is not scalable enough to accommodate all pile arrangements a user might create. For example, several piles can be made to occupy the top edge of the PDA, because pile markers are very space efficient. Our current history browser would be unable to accommodate such density. Alternatives are in the realm of future work.

#### SHARING SCENARIO 1: WIRELESS P2P OPERATION

We now turn to scenarios where a small group of co-located PDA owners wish to operate on piles together. For example, consider a returned traveler reporting on a trip while showing photographs. Traditionally, the story teller comments on each photo as he starts passing it around. A stream of images makes its way around the circle of listeners as the story unfolds. The further down the chain, the further out of phase with the story a listener is once a photo arrives. This problem could be remedied through a single large display that everyone views. But this solution has shortcomings too. One is that such a central gathering place is often not available when the mood strikes to show photos; a restaurant venue, for example. Another is that listeners often have differing interests in any given photo. Some may, for example, wish to linger over the image of a friend who is unknown to the others.

We try to provide a better solution by enabling PAS to share virtual piles among PDAs. This facility allows participants in the above scenario to look at the photos at their own pace. Beyond this advantage, our solution provides a further benefit for the photo story scenario and analogous ones in professional settings: participants can impose their own organization on the photos by creating their own local piles. They might, for example, create a pile of shots that they'd like printed. Figure 3 shows how two users interact through the peer-to-peer (P2P) facility of our system. They are sharing two piles, symbolized by the circles with the ocean photo and the palm tree sunset. The current implementation uses ad hoc wireless (WiFi) networking for inter-PDA communication.

Note that a given pile is not necessarily located at the same place along the edge of each PDA's screen. For example, for

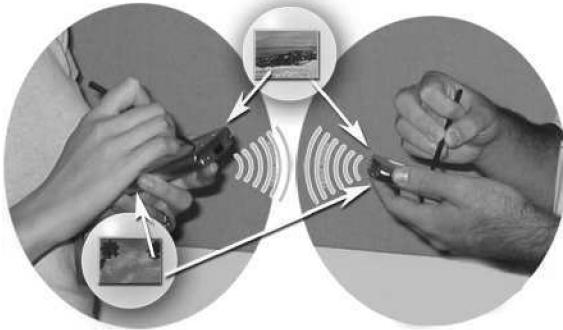


Figure 3: Two PDAs share two virtual piles.

the participant on the right, the ocean pile is logically located at his upper right, while the same pile is located at the upper left for the participant on the left side of the Figure. The model is for the participants to agree on a pile position in space by pointing out to each other verbally where they intend to place the pile. This way they can then communicate about the information during subsequent verbal interactions. An intuitive use of PAS is thus for each participant to place a pile marker such that it references its pile's imagined place in space. The system does not, however, enforce such alignment. At least in the current implementation there is no sensor that would detect where participants point as they agree on a pile.

Pile sharing is initiated by the person who has created the pile. The user activates sharing from a context menu (stylus point-and-hold). This action causes a popup dialog to open on all PDAs within wireless 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, each recipient user is free to place the newly shared pile anywhere along the edge of his screen by dragging an automatically created pile marker into position. From then on the user can interact with the shared pile by dragging it into the work area, or by flicking new information to the pile marker. Once a pile is shared with a set of participants, all participants can thus add and remove items from the shared pile. Each such action is broadcast and executed by all participants' PDAs.

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; there is no notion of serializability and concurrency control, since the user scenario is of a few people working in reasonably close proximity. We do not envision thousands of PDAs sharing 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. 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



Figure 4: Three participants by the table display.

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, pile sharing is discontinued. 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.

#### SHARING SCENARIO 2: TABLETOP DISPLAY

Beyond the P2P capability we described above, we integrated our Piles Across Space implementation into the environment of a DiamondTouch tabletop display. This apparatus is a (diagonally measured) 107cm table whose surface is a 1280x1024 pixel touch-sensitive, top-projected display. The table is more than a large, horizontal display surface in that it allows up to four persons to touch the surface of the display at the same time, and to drag items across the table's expanse. Through capacitive coupling across the pads that participants sit on, the attached computer can distinguish touch from all hands, even when they touch the surface simultaneously.

We use this table in the context of our research on using PAS as a tool for field biologists. One scenario is of students and professional biology researchers returning to their field station from a day outdoors. They are seeking advice from experts, or simply want to share photos or interesting measurements with others. Our vision is for them to flick piles of information from their PDA to the table for others to see, manipulate, and return to the PDA.

Figures 4 and 5 show an overview and a screenshot of the tabletop arrangement. Figure 4 shows two users seated, and one standing with a PDA. The PDA communicates with the table computer through ad hoc WiFi. The interaction model between the two devices is more loosely coupled than our inter-PDA P2P approach. Information is simply transferred between the two devices, rather than continuously maintained through shared piles.

The light colored area on the table (near the arrows in Fig-

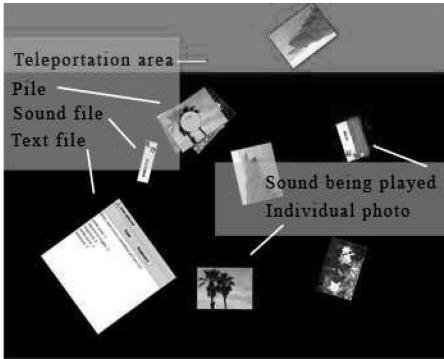


Figure 5: Screenshot of the tabletop display with annotations superimposed.

ure 4) is called the *teleportation area*. The PDA’s screen in Figure 1 shows the word ‘Table’ along the left edge. This is the small device’s equivalent to the table’s teleportation area. To transfer either a pile or an individual information item from the PDA to the table, one drags it to PDA’s teleportation area. As soon as the respective icon is dropped, it appears in the table’s teleportation area. Anyone seated at the table can then drag the item freely around the surface of the table display.

Figure 5, a screen shot of the table surface, shows a number of items strewn across the table. All came from PDAs running PAS. As the superimposed annotations show, the lower right item is a text file. Above it we see a sound file. An entire pile lies above the sound file. Piles are a first-class data structure in our table application. Piles can be modified, created, or destroyed by groups of users at the table, and then transferred back to the PDAs. At the right edge of the table we see a sound file being played; an audio control has popped out of the main window. The rest of the items on the table are individual photos that were flicked to the PDA’s teleportation area.

Conversely, piles or individual information items that table operators drag to the table’s teleportation area accumulate as a special pile on every PDA that is in range. The presence of new information is indicated on the PDA by having the lower portion of the PDA’s teleportation area highlighted. The device’s owner can then drag it onto the work area.

## RELATED WORK

Our non-intrusive pile markers are similar in appearance to Baudisch & Rosenholtz’s [2] “halos” though the functionality differs, since our markers are not intended to convey off-screen distances through variations in radius size. Other systems use arrows that point off screen instead, such as [7] and [4]. We prefer halo-like markers to arrows, since the former consume less screen real estate.

The general concept of storing information off screen is also found in earlier 3D desktop technologies such as Toolspaces [11]. Similarly, [8] is a recent, small screen oriented focus + context reference.

Several research projects explore combinations of individually-owned devices with large, shared displays. The Pebbles

project [10] explores the use of PDAs to control applications on a shared vertical display. The STARS project [9] uses PDAs in combination with a tabletop display to show secret game information. The UbiTable [12] allows pairs of users with personal laptops to transfer information between their own devices and a shared tabletop display. Piles are not the central data structure in that work.

## CONCLUSION

We described Multi-User Piles Across Space. The original PAS application allowed a single user with a PDA to manage off-screen piles of information. We extended this technique to support co-located operation, either among multiple PDAs or between PDAs and a shared tabletop display. Multi-User Piles Across Space allows users to share the contents of virtual piles and to work together to organize information even given the screen real estate constraints of PDAs. When a large public display is available, we allow users to take advantage of it, while still retaining piles as first-class citizens in order to support information organization activities.

## REFERENCES

1. Open Mobile Alliance. Wireless application protocol. [http://www.openmobilealliance.org/tech/affiliates/wap/wapindex.html#wap20\\_2001](http://www.openmobilealliance.org/tech/affiliates/wap/wapindex.html#wap20_2001).
2. P. Baudisch and R. Rosenholtz. Halo: a technique for visualizing off-screen location. In *CHI’03*, pages 418–488, 2003.
3. O. Buyukkokten, H. Garcia-Molina, and A. Paepcke. Seeing the whole in parts: Text summarization for web browsing on handheld devices. In *10th International WWW Conference*, 2000.
4. D. Cox, J. Chugh, C. Gutwin, and S. Greenberg. The usability of transparent overview layers. In *CHI’98*, pages 301–302, 1998.
5. O. de Bruijn, R. Spence, and M. Y. Chong. Rsvp browser: Web browsing on small screen devices. *Personal Ubiquitous Comput.*, 6(4):245–252, 2002.
6. G. W. Fitzmaurice. Situated information spaces and spatially aware palmtop computers. *Communications of the ACM*, 36(7):38–49, Jul 1993.
7. B.L. Harrison, G. Kurtenbach, and K.J. Vicente. An experimental evaluation of transparent user interface tools and information content evaluation. In *UIST’95*, pages 81–90, 1995.
8. E. Lank and S. Phan. Focus+context sketching on a pocket pc. In *CHI ’04*, pages 1275–1278, 2004.
9. C. Magerkurth, M. Memisoglu, R. Stenzel, and N. Streitz. Towards the next generation of tabletop gaming experiences. In *Proc. of Graphics Interface 2004*, pages 73–80, 2004.
10. B. Myers, H. Stiel, and R. Gargiulo. Collaboration using multiple pdas connected to a pc. In *CSCW’98*, pages 285–294, 1998.
11. J. S. Pierce, M. Conway, M. van Dantzich, and G. Robertson. Toolspaces and glances: storing, accessing, and retrieving objects in 3d desktop applications. In *SI3D ’99: Proc. of the 1999 symposium on Interactive 3D graphics*, pages 163–168, 1999.
12. C. Shen, K. Everitt, and K. Ryall. Ubitable: Impromptu face-to-face collaboration on horizontal interactive surfaces. In *Proc. of UbiComp 2003*, pages 281–288, 2003.
13. QianYing Wang, Tony Hsieh, and Andreas Paepcke. Piles across space, breaking the real-estate barrier on PDAs. Technical Report 2005-17, Stanford University, July 2005. Available at <http://dbpubs.stanford.edu/pub/2005-17>.
14. K. Yee, K. Swearingen, K. Li, and M. Hearst. Faceted metadata for image search and browsing. In *CHI’03*, pages 401–408, 2003.