Collaborative Personal Information Management with Shared, Interactive Tabletops

Anthony Collins
School of Information Technologies
University of Sydney, Australia
anthony@it.usyd.edu.au

Judy Kay
School of Information Technologies
University of Sydney, Australia
judy@it.usyd.edu.au

ABSTRACT

Given the central and routine nature of Personal Information Management (PIM) to conducting daily tasks, the ability to conduct PIM with tabletop interfaces (and other collaborative devices that diverge from traditional desktop interaction) is an important method of supporting collaborative work at the tabletop. Using tabletops potentially allows people to conduct PIM in a comfortable place that is used for a range of regular activities, from eating to reading and writing. However, there are challenges in supporting natural and useful PIM with the properties and limitations of tabletop interaction. The use of large surfaces also provides great possibilities for collaboration, as PIM often involves other people, or people may want to selectively share their personal information with others for a variety of purposes (such as organising work e-mail with colleagues).

In this paper, we present our techniques for navigating and sorting multiple sets of personal information—particularly digital files and e-mail—on a tabletop. We present *Focus*, a multi-user tabletop interface where users select "focus items" that retrieve related artefacts from unified personal information spaces, with minimal prior configuration. We conclude with a discussion of results obtained in a small-scale user study, where participants were asked to access and organise personal information at an interactive tabletop.

Author Keywords

Personal Information Management, Tabletop Interface, Single Display Groupware

ACM Classification Keywords

I.3.6 Computer Graphics: Methodology and Techniques—*interaction techniques*; D.2.2 Software Engineering: Design Tools and Techniques—*user interfaces*; H.5.3 Information Interfaces and Presentation: Group and Organization Interfaces

CHI 2008, April 5 - 10, 2008, Florence, Italy.

INTRODUCTION

Personal Information Management (PIM) is a task that is typically performed with a conventional personal computer, by a single user. Moving away from the desktop paradigm and its single-user interaction methods allows new possibilities for both sharing personal information, and collaborating on the management process. The interactive tabletop—a novel medium that has recently attracted significant research interest—supports collocated collaboration with a large, shared workspace, where people can sit face-to-face and interact with it simultaneously.

While PIM is typically a personal task, we would like to support people naturally sharing and accessing their personal information at a shared, interactive tabletop. A key facility of a tabletop is to support interaction with digital files for conducting collaborative tasks. Considering the central nature of other types of personal information (such as e-mail) to conducting daily tasks, support for accessing and organising this information at the tabletop can potentially be very useful. For example, an office-worker may want to share work related e-mails and files with a colleague in order to perform decision making. Another example is two people working on a joint project involving different types of information, where they have divided the work. The pair would like to combine their work and discuss it using a collaborative interface.

Despite the advantages of the tabletop medium for collocated collaboration, the properties and constraints of tabletop interaction call for re-thinking standard approaches to accessing and managing personal information. First, it must be natural to interact with the tabletop using input with special constraints not present in personal computer interaction: a keyboard and mouse is typically not present in a multi-user tabletop setting. While providing a projected keyboard on the tabletop surface is possible, this provides no tactile feedback when pressing keys, and does not give users a fixed reference of where to place their hands.

Second, tabletop interfaces must support people sitting in a range of orientations, as a fundamental facility is to support people sitting both face-to-face and around-the-table. Consequently, interface elements must be orientation independent. This makes the use of text (such as file and directory names) on a tabletop problematic, as it may be difficult to read text from all positions around the tabletop.

Third, the size of the tabletop has an impact on a number of design factors, such as the size at which to display text and interface elements. More importantly, the limitations of physical human reach and direct-touch interaction must be considered, as a user may not be able to reach across the table. User interface selection targets must also be sufficiently large for direct-touch interaction, which potentially contributes to clutter on the tabletop. A tabletop PIM interface must support users working with a particular set of personal information at any position around the tabletop—this means that the point of interaction with a personal information set must not be dependent on a particular position (or orientation) on the tabletop surface.

Finally, when considering tabletop size, display resolution, and multiple users, clutter is a critical problem to address in tabletop interfaces. Given the increasingly large size of personal information collections, the need to present large amounts of information, or the need to replicate information on the tabletop for viewing or interaction by multiple people in different orientations, is a major problem when designing to minimise clutter. Multiple people may create and interact with artefacts that overlap both shared and individual objects on the tabletop, particularly when artefacts are zoomed-in for a more detailed view. In a multi-user setting, where several people are creating and interacting with content on the tabletop simultaneously, management of clutter becomes challenging—existing methods to manage clutter in conventional personal computers, such as a 'task-bar' in a window manager to control active windows, are only designed for single-user interaction. Another challenge is that the management of clutter by one user could have unintentional and adverse affects on other users of the tabletop.

When taking into account the properties of tabletops, and their collaborative nature, we need to explore ways to access and organise digital information in a natural way at the tabletop. The unified approach to personal information management is ideal for conducting tasks at the tabletop, as it moves away from users navigating separate *silos* of information at the tabletop (such as uniquely organised hierarchies for file systems and e-mail). In this paper, we explore multi-user, unified access to e-mail, personal files, photos, and stored web-pages using interactive tabletops.

RELATED WORK

A significant problem in modern personal information management is information fragmentation [16], where people regularly need to access and manage related information in separate physical locations with little support from the tools they use. Studies conducted by Boardman and Sasse [4] highlighted that people employ a range of personal strategies to manage their information, both within and across different PIM tools. The notion of *folder overlap* was observed for many participants, particularly for files and e-mail, where people create a similar hierarchical organisation between the two information spaces. As we are exploring collaborative PIM using a tabletop, we must support people concurrently interacting with information that is organised differently, both within a person's information collection, and between multiple people's information collections.

The *Stuff I've Seen (SIS)* interface provides a unified index to all files and personal information that a user has seen on their computer [8]. SIS does not rely on a hierarchical organisation structure, and uses contextual cues (such as dates and people) to facilitate information retrieval. This is similar to *Lifestreams* [12], and *MyLifeBits* [13], where the information a user has seen is ordered on a timeline for later retrieval. Improved search facilities for personal information, which merge searching and browsing, were presented with *Phlat* [7]. Unified tagging allows users to create an organisation scheme consistent across all of their personal information. Furthermore, Phlat (like SIS) presents this information in a single interface, thereby removing the need to use separate tools to access different types of information. These projects tackle the problem of supporting a unified personal space of information.

Content-based search, such as with *Google Desktop Search*[1], *Apple Spotlight*[2], and *Windows Vista Instant Search*[3], also allows interaction with a single information-space that merges relevant content regardless of its type and where it is stored. However, these content-based access mechanisms provide a targeted search interface that would be difficult to adapt to a collaborative tabletop, where there is a need to support interaction with multiple personal collections that may often contain related content.

Haystack [15] is a general-purpose information management tool that allows a variety of information types to be managed within a single user interface. Users can also create arbitrary relationships between the stored information to simplify retrieval. Karger *et al.* (*ibid.*) discuss the concept of presenting "similar items" alongside user-selected items, in order to support iterative query refinement, which is similar to ideas explored in the *IQ prototype* [9], where information relevant to the the user's current context is presented in a peripheral interface.

The notion of *Group Information Management (GIM)* has been a recent topic of discussion [10]. The focus of GIM is on sharing information—that was created with a clear intent to share—with groups and institutions, rather than individuals in a collocated context. Our focus in this research is more accurately viewed as collaborative PIM (rather than GIM) where people work collaboratively with their personal information, or they selectively share parts of their personal information with a small number of people in a collocated context.

PIM research has focused on making organisation and retrieval easier, but there has been little exploration of conducting PIM with devices that are inherently collaborative. These devices need to support people sharing their personal information with others, as well as helping people collaboratively navigate multiple sets of information.

[1] `http://desktop.google.com/`
[2] `http://developer.apple.com/macosx/spotlight.html`
[3] `http://www.microsoft.com/vista/`

Tabletop and pen-based interface research has explored interaction with small collections of information, such as digital photographs [3], or a virtual desktop of files [1]. Pure hierarchical interaction has been explored in the *Personal Digital Historian (PDH)* project [23], as a way of increasing the scalability of a tabletop interface. Interaction through passive and active associations has also been explored [22], although results suggest that participants highlighted "clutter" and "over crowding" within groups of information as major problems in the interface. This crowding issue is noted as an important factor to address in tabletop interfaces [19].

Research into PIM at a tabletop has been limited. Earlier work on *micro-mobility* [11] has researched people explicitly and visibly transferring files between devices, such as tabletops and personal computers. Other approaches require users to move their information into a shared region of the tabletop from a laptop or personal storage device [14, 21]. These systems do not support people interacting with broader collections of personal information at a tabletop.

Techniques for collaborative searching of large collections on tabletops have been explored in *TeamSearch* [17], although these require manual Boolean query formation. However, the benefits of collaborative searching are highlighted, such as facilitating stronger collaboration and awareness among group members. Providing a search interface alone is not sufficient for supporting natural browsing of multiple sets of personal information on a tabletop.

We previously explored *OnTop* [6] as an interface for collaboratively accessing personal file systems from a tabletop. Our evaluation showed that OnTop provided valuable support for tabletop collaboration on the *access* process—participants were more social and co-operative when using OnTop instead of a hierarchical alternative (which encouraged private file system access and selective sharing of files once they had been located). As OnTop presents information according to its content, rather than storage location in a file hierarchy, the techniques used are applicable to a variety of organisation schemes and types of personal information.

BRINGING PIM TO THE TABLETOP

Previous PIM research has explored different techniques to support people interacting with their personal information [7, 8, 12, 13, 15, 18]. Support for people interacting with group information is also an active research area [10]. However, there has been little research into supporting people conducting PIM using collaborative, multi-user devices, such as tabletop displays. Despite this, there is a need to support PIM at the tabletop in order to facilitate people collaborating with personal information in a collocated setting.

The following scenario illustrates collaborative PIM using a tabletop interface:

> Jim is eating breakfast while sitting at his dining table, which is also an interactive tabletop. As his wife, Sarah, is also sitting at the table, they wish to make some decisions about an upcoming holiday trip. They have been organising important information using e-mail, text documents, photos, and stored web-pages, on each of their personal computers. They must make some bookings during the day, and they have decided to split the workload. Jim needs to access his e-mail (stored on his desktop computer in another room), which contains quotes for hotel bookings, and needs to exchange this with Sarah who is organising transport to and from the hotel. Sarah needs to exchange details about when she is able to take time off work for the holiday, with the information stored on Sarah's desktop computer (also in another room). They use the tabletop to jointly review e-mails, files, photos and web-pages, in order to make important decisions about the trip, and to delegate different aspects of the planning.

This scenario involves sharing of personal information that needs to come from multiple remote computers. Due to the context of use of tabletops (which will often be used in shared, open environments), we must support people interacting with PIM collections that are stored remotely—we do not envisage that a tabletop is going to be used as a desktop computer replacement. Rather, tabletops and other devices (such as desktop and laptop computers) must operate together harmoniously. This requires universal access to personal information and its associated meta-data from a variety of devices in a pervasive computing environment. We focus on providing support for collaborative PIM using a tabletop.

When supporting multiple people sharing their personal information, we need to facilitate interaction that is independent of both the PIM tools used and the underlying storage locations of the information. Different tools (such as an e-mail client, and a file system interface) may organise related information in different hierarchies, with little integration between them. When accessing this information at the tabletop, we need to present a unified view of this information that is independent of its storage. If people bring their personal information to the tabletop, we need to facilitate natural interaction with information from a variety of tools—if one person is using a certain e-mail client to store their mail in a folder hierarchy, and another person is using a different client that uses tagging instead of hierarchical organisation, we need to support both these people easily interacting with the mail in a consistent way at the tabletop.

Furthermore, in the scenario presented, the two sets of information that are accessed at the tabletop contain highly related information. We need to support collaboration with the combined set of information at the tabletop. As a key goal of our work is to support people collaborating with each other, it is likely that people will be sharing related information, and so we envisage that there may be significant information overlap between collections. Therefore, we need to support people interacting with similar collections of information.

With these design considerations in mind, we now outline our techniques for supporting collaborative PIM at a shared, interactive tabletop.

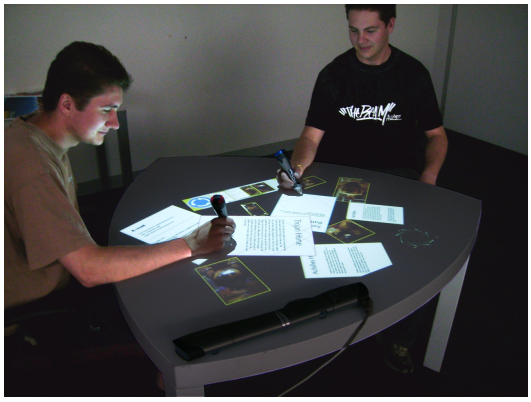

Figure 1. People using Focus to collaboratively access and organise their personal information at a tabletop.

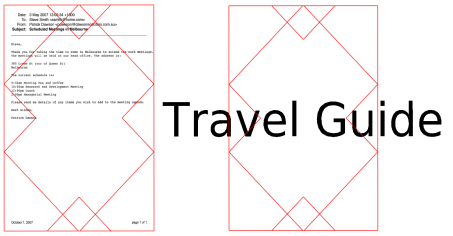

Figure 2. A selected e-mail (left), and a flipped-over PDF document (right) with its filename written on the back.

FOCUS

To provide the functionality discussed in the previous section, and to address the properties and constraints of tabletop interfaces, we have created the *Focus* interface for associative access to multiple sets of personal information at an interactive tabletop. Focus (shown in Figure 1) presents a unified view of personal information, that is independent of the physical storage location of personal information, and the tools used to manage it. Focus provides a *remote* interface to personal information, where users *export* selected personal information to the tabletop for collaboration.

User View

Focus shows image representations of personal information that appear to be placed on top of the tabletop surface. E-mails, PDF documents, photos and web-pages are presented and manipulated on the tabletop in an identical way. The appearance of e-mail on the tabletop is similar to a printed e-mail, with a brief header describing the mail recipients and subject. An artefact can be *moved* by selecting it in the main area inside the solid selection lines (shown in Figure 2) and dragging it. When moving objects, they have realistic momentum so they can be *flicked* around the tabletop, making it easy to move objects to areas of the tabletop that are out of physical reach. Objects on the tabletop can be rotated and resized (in a combined *rosize* gesture) by selecting the object at one of its corners and dragging the corner.

A personal information object can be *flipped* by selecting it from within a stippled triangle along one of the its edges and dragging it across to the opposite edge. Once flipped, the user sees a description of the object written on the back, as shown in Figure 2. For example, if the object is a file, then the filename is written on the back.

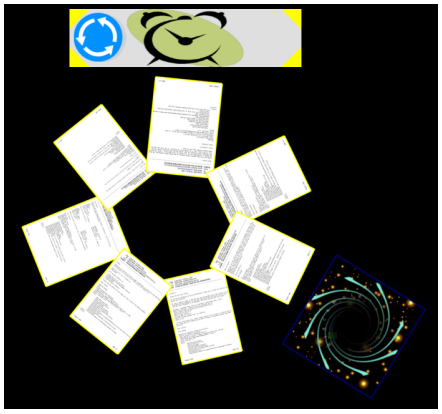

Figure 3. The Focus start view, which shows the contents of a user specified mailbox in a radial layout on the tabletop.

When Focus is first launched, a broad *start view* of the personal information is presented in a radial layout on the tabletop (shown in Figure 3). Our current implementation has two types of start views available. The first shows the first file (alphabetically) in each exported directory of each remote file system, and the most recent e-mail in each exported mailbox. This should work well if users have a reasonable organisation of their files, and if the first file in a directory (or the most recent e-mail in a mailbox) is representative of the folder or mailbox content. The second technique presents only the contents of a nominated mailbox on the tabletop (in the same radial layout). From these start views, the content presented on the tabletop can be used to navigate to other related information contained in the personal information sets. In this paper, we present and evaluate the mailbox start view.

Navigation of personal information is based on the notion of a *focus* item. Once a user selects a focus item, all other *related* pieces of information are displayed on the tabletop (regardless of where they are stored or which computer system and user they belong to). The personal information from multiple people is merged to appear as a single, combined information set on the tabletop, and a user may navigate the information by re-selecting focus items. To select a focus item, a user dwells (depresses the pen for one second) on it. A 'click' sound gives feedback that it has been selected. Matching results from each remote computer are returned and displayed immediately on the tabletop. This approach allows people to easily find related information in both familiar and unfamiliar collections.

Figure 4 shows Focus after two focus item selections from the start view. To address the limited display area of the tabletop, personal information is presented in a non-uniform manner, where the initial size of each items' image representation is determined by its relevance to the focus item. Thus, most relevant items appear large and prominent, while

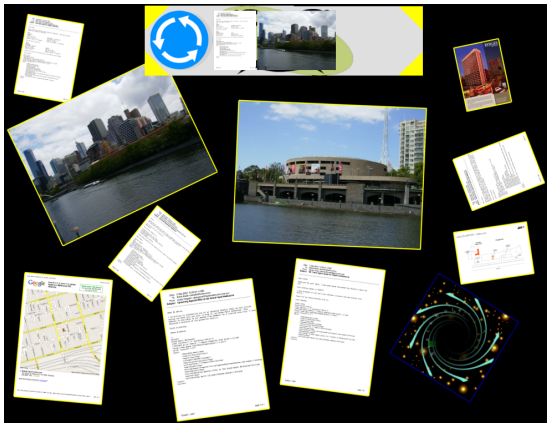

Figure 4. The Focus interface showing a collection of e-mails, photos, PDF documents and stored web-pages, that have been rearranged by users after two focus item selections.

less relevant items appear small and unobtrusive. The resize gesture can be used to override the initial size based on similarity. If an item was previously displayed on the tabletop, it is presented in its last location if the user moved it to a specific position. This is to provide spatial consistency between focus selections, and to allow users to create arbitrary spatial groupings of their personal information (as in *Data Mountain* [18]).

To reduce clutter on the tabletop, users can move unwanted items in the *Black Hole* (shown in the bottom right of Figures 3 and 4), a special object that is always visible and above other objects. It can be moved, rotated, and resized like any other object, although an object reduces in size as it is dragged closer to the centre of the Black Hole until it is completely hidden. Objects placed in the Black Hole may be retrieved by reducing the hole's size and flicking it with sufficient momentum, causing the contained objects to "fall out".

The *History Browser* widget, shown at the top of Figures 3 and 4, supports 'back' and 'forward' navigation operations. This object can be manipulated like any other interface object, although it cannot be placed in the Black Hole. It shows thumbnail representations of recent focus items (typically the last five focus items, depending on the aspect ratio of the thumbnails), and a user dwells on a past focus item to go 'back' to the display of information associated with it. A special circular icon on the far left of the History Browser represents the original start view, which can be returned to by dwelling on the icon.

Storage Bins, shown in Figure 5, are special folder-like objects on the tabletop that show thumbnails of items dragged into them (similar to [20]). We use a space-filling thumbnail layout to accommodate the thumbnails placed in the Storage Bin. Storage Bins can be moved, rotated and resized on the tabletop just like any other interface object, and they can also be placed in the Black Hole.

Figure 5. An empty Storage Bin (left) with a pre-defined label, and a Storage Bin containing four items (right) with its label obscured. Items can be removed from the Storage Bin by dragging them out.

Exporting Personal Information

The Focus interface supports people interacting with personal information stored on a remote computer. As a result, Focus is designed so that people export a subset of their personal information to the tabletop when they need to use it in a collaborative setting. To make personal information accessible by the Focus interface, users run the *Focus Exporter* on their personal computer. After a focus selection from the tabletop, a query is sent to each connected machine to find similar personal information. This query request contains meta-data associated with the focus, so that similar information can be found in separate personal information spaces that do not contain the original object.

A critical issue in supporting this functionality at a collaborative tabletop is privacy. To address this, users are given explicit control over the information accessible by tabletops—users may specify (using the Focus Exporter application) which tabletops may connect to the computer to request similar items, and configure which information can be returned. Currently, the tabletop access control is based on IP addresses, as it is expected that all the computer systems are connected to the same local trusted network. Users may also configure the start view of their information. This involves either selecting a file system path where the first file in each directory, and the most recent e-mail in each mailbox is presented for the start view, or selecting a mailbox to be displayed on the tabletop. As the start view is configured on each remote machine, each person can have their own customised view presented at the tabletop.

Similarity of Personal Information

As navigation with Focus is performed by association, where similar information is presented after a focus item is selected, it is essential that Focus will correctly retrieve the relevant information. It would be unreasonable to implement a static relevance calculation mechanism, as people use a variety of tools to mange their personal information, and they employ a variety of different organisation strategies both within and across different tools [4]. Accordingly, we have created a mechanism for users to customise how their information is determined relevant to a given focus—users select (on a scale of 0 to 10) how important certain meta-data attributes are to the organisation of their information, as shown in Figure 6. For example, if the user deems information by the same author(s) (such as the document author, or the sender of an e-mail) as closely related, the user increases the weighting on the 'authors' attribute. A 'tool-tip' is displayed when moving the mouse over each meta-data attribute to describe its meaning. For example, moving the

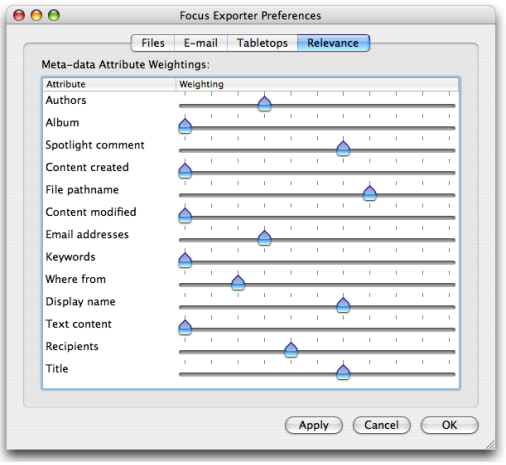

Figure 6. The Focus Exporter application allows a user to personalise the relevance calculation of their information. The weighting of each meta-data attribute may be adjusted based on the content and organisation of the user's information.

mouse cursor over the "Content created" attribute will pop-up the text: "Items that were created on the same day and within the same hour are deemed relevant to each other". The relevance customisation is performed on each remote computer to suit the semantics and organisation of the information.

Currently, the weighting of 13 meta-data attributes can be adjusted, including the full text content of information that contains text (such as e-mail and PDF files). The attributes given a high weighting by default are keywords, file pathname (for items in the same directory, or e-mails in the same mailbox), display name (such as the filename), title (such as an e-mail subject line), and text content.

Implementation

Focus has been developed using the *Cruiser* tabletop platform [2], a multi-user, gestural, collaborative tabletop system. It uses the *Mimio Capture* whiteboard pen system, but is independent of the hardware (a version called *SharePic* [3] used the *DiamondTouch*). Cruiser uses multi-threaded C++ with OpenGL and leverages capabilities of modern graphics hardware and multi-core processors.

Focus currently supports a range of information types, including images (of most popular image types, such as JPEG and PNG), PDF documents, stored web-pages, and e-mail from any e-mail client that supports the *mbox* or *Maildir* formats. Currently, the first page of a PDF document, web-page or e-mail is presented in the interface.

Focus Exporter is a separate application written using the Apple Cocoa and Spotlight search frameworks. Because an existing desktop search framework has been used, which treats information from different PIM tools (including the hierarchical file system) in a uniform manner, it scales well to large collections of personal information. The difference in search time between a small collection of information and a large collection is unnoticeable. With the Focus Exporter running on an average desktop computer, results start to be displayed on the table within 2 seconds of a focus selection. Image representations of items are loaded on the table while the computer is retrieving further results in the background, so that users immediately see the effect of their focus selections. Accessing multiple sets of personal information, rather than a single set, has no affect on the overall time taken for relevant items to be presented on the table, as the retrieval process is parallel and asynchronous.

In our current implementation of Focus, a list of remote computers to connect to must be specified before launching the tabletop interface. In future, we will use multicast DNS service discovery to automatically detect exported personal information on a local network (with privacy concerns in mind).

EVALUATION

We conducted a qualitative evaluation to assess the appropriateness of presenting a unified view of personal information on a tabletop, and whether users are able to effectively navigate their personal information by selecting focus items. Our goal was to examine how people use the e-mail start view of Focus (where the contents of a mailbox is initially presented) as a way of finding and organising their personal information at the tabletop. At this stage, our evaluation was a single-participant study, as we wanted to analyse the usability with one participant interacting with the Focus interface. In future, we will conduct evaluations with groups of participants interacting with multiple sets of personal information at the tabletop, as a key goal of Focus is to support collaborative PIM.

While the long-term field study is a popular evaluation technique for PIM due to its complex and ongoing nature [5, 7, 8], we needed to first examine the usability issues in supporting PIM with a novel tabletop interface. Consequently, our evaluation was designed to gain qualitative usage data by having participants complete a set list of tasks, with fictitious personal information (which users were given sufficient time to familiarise themselves with). The "personal information set" was based on the real personal information of one of the authors, making it a realistic collection.

Six participants were recruited for the evaluation, from a range of backgrounds and experience levels. Most were faculty members and students, though they all had limited experience using a tabletop interface. The ages of participants ranged from 22 to 35.

User Tasks

Participants worked individually on three PIM tasks that were designed to be meaningful to the recruited participants. The tasks were related to the travel scenario previously outlined, where the participant had to organise information related to the tasks that they were going to complete during the day, based on the contents of their e-mail inbox, and various files (such as web-pages, PDF documents, and photos) stored on their computer. Each task involved retrieving items related to a particular group of e-mails, and placing them in

Storage Bins. A Storage Bin was allocated for each task, with the title on the Storage Bin indicating which task it was for.

A pre-organised information set was provided for participants to use in the experiment. This ensured comparability between experiments, and it avoided potential privacy issues. The data set contained 10 e-mails, and 53 files (PDF documents, images, and stored web-pages), sorted into 6 directories. Participants were given a printed list of the organisation of e-mails and files that they were free to refer to at any time during the experiment (although the tasks did not require a knowledge of where the information was physically stored).

The number of e-mails and files used in this experiment was a plausible number for the scenario given to participants. As the context of use of a tabletop differs from that of a conventional personal computer, we expect that users will only export information to the tabletop that they need to collaborate with, or need to access at the tabletop for a specific purpose. The number of files and e-mails was also enough to clutter the tabletop, meaning that it would be unworkable if all were presented at once. To be more realistic, the information set also contained some information relevant to the scenario, but not relevant to the user's tasks.

Keywords were assigned to the information in order to enhance the relevance calculation performed by Focus, particularly for images as the content-based search could not be fully utilised. The meta-data weightings for relevance calculation were fixed between experiments, with a high weighting given to textual document/e-mail content, keywords, filename, and the e-mail subject field.

Participants were asked to complete the following tasks[4] at their own pace:

1. You have booked a flight, but there is a mistake on the booking confirmation. Find the booking information and place it in the indicated Storage Bin so you can later rectify the problem.
2. You need to book a hotel at your destination. Find the information you have gathered about suitable hotels, and choose which hotel you would like to stay at. Once you have decided, place the relevant information in the indicated Storage Bin so you can later make the booking.
3. You need to book a rental car to use during your stay. Find the information you have gathered about the cheapest cars available. Once you have decided on a car, place the relevant information in the indicated Storage Bin so you can later make the booking.

Experiment Procedure

The experiment began with the experimenter introducing the activities. The participant was given a brief tutorial on the Cruiser interface—core functionality of the interface was emphasised, such as rotating, resizing and moving objects, and using the Black Hole to hide unwanted objects. After the introduction, the participant was given an instruction sheet to work through at their own pace. It began with basic tasks, such as moving objects around the tabletop and placing them in Storage Bins, and moved on to tasks involving focus selections. This introductory part of the experiment was designed to be completed in ten minutes, and the experiment did not proceed until the participant understood the concepts and interaction techniques introduced to them.

After the participant completed the first section of the instruction sheet, the interface was reset to a default state before completing the task list outlined in the previous section. An e-mail inbox was presented in a radial layout for the start view, with Storage Bins for each of the tasks positioned around the perimeter of the tabletop. The participant was then instructed to complete the three tasks specified on the instruction sheet.

After completing the tasks, the participant was asked to answer a short questionnaire about their background, and about their experiences using the interface. The evaluation was captured on video, and the experimenter later analysed this to identify interesting features of interaction.

RESULTS

All participants were able to complete the tasks without any significant obstacles. We observed participants employing a variety of techniques for navigating the personal information on the tabletop. E-mails were used extensively as focus items by all participants to navigate through the information, even though in many cases they could have selected other items instead (such as photos and PDF documents). All participants could navigate the information easily by using the e-mail start view, as the e-mails contained in it were highly representative of the information required for the experiment tasks.

Two participants employed an iterative query refinement strategy, where they continued making more specific focus selections to reduce the number of objets displayed on the tabletop. The remaining four participants simply navigated to the broadest focus item that was related to the information they were required to find (meaning that there were many, often loosely related items on the tabletop). Despite the additional clutter, no participant encountered any difficulties managing the information on the tabletop, as they were quick to identify items related to their task.

For all participants, digital photographs, web-pages and PDF documents that contained pictures and business logos were faster to identify on the tabletop (as participants always interacted with them first), and so they were used as focus items more often (unless the start view was presented, where only e-mails were present on the tabletop). However, it is important to note that the issue of identifying graphical representations of personal information would be strongly affected by the participant's familiarity with the information.

[4]The tasks in the list have been shortened—the actual task list contained specific details about the information that participants were required to find.

Surprisingly, participants rarely used the History Browser to return to past focus items or the less cluttered e-mail start view. The start view contained e-mails that would provide easy access to the items required for each task, though it was used as a last resort by all participants (when the participant could not find anything related to what they needed on the tabletop), with one participant stating that it would be "inefficient to use it when they can focus on something already on the table instead."

As the interface retains the position of objects on the tabletop between focus selections, all participants used a grouping strategy on the tabletop. For example, if the participant was completing a task where they needed to find hotel details and they encountered information about rental cars (required for a later task), they would move the rental car information to a separate area on tabletop. However, the same positions were sometimes used for different objects when another item was focussed on. This often caused the objects on the tabletop to overlap. Despite this, only one participant noted this as a difficulty in using the interface—we observed that even though objects overlapped, participants examined the content of each object as they moved it out of the way. The use of animation when information was loaded onto the tabletop made it clearer that objects were being loaded to the same location (where they were overlapping).

As Focus is designed to support people sitting at any position around the tabletop, it initially presents information on the tabletop in a radial layout. Participants were free to sit at any position around the table during the experiment, though no participant decided to change positions during the tasks. This meant that participants needed to regularly reorient objects on the tabletop so they could view them properly from their chosen position.

A key issue emerging from the questionnaire responses was personal information privacy. Participants were asked if they would be comfortable using the Focus interface to share their personal information at the tabletop with a friend or colleague. Four participants explained that if they were to use such an interface, they would need complete control over the search scope. One participant stated that they would only use Focus to connect to a machine that did not store private and sensitive information, but would be comfortable using it with a work machine.

Many usability suggestions were made by participants in the questionnaire. For example, one participant suggested that each thumbnail in the History Browser should have a caption (such as the filename), or some kind of identification mark that helps to distinguish between thumbnails that look similar (such as a group of e-mails containing only text). One participant suggested that the system remember when they place items in the Black Hole (but only remember this for the current focus selection). This means that the Focus relevance calculation mechanism should effectively learn what is irrelevant to a given focus item.

When asked about being presented with a unified view of personal information on the tabletop, one participant stated that it "felt natural having all of the different items appear on the table—it made finding things related to a topic easier." Another participant stated that Focus was "easier to use than a conventional interface—especially for things that are stored in different ways (e.g. e-mail and photos)." One participant commented that "Focus integrates otherwise separately kept information, which is needed for tasks, and I like that." Another participant noted that it was "trickier when there was more noise in the returned information" (related to the Focus relevance calculation mechanism, which returned some results not completely relevant to the task the user was completing).

Finally, participants were asked about their current personal information storage system. No participant said that they sorted their e-mail collection in the same way as their file system. Four participants stated that they did not have a predetermined categorisation scheme for their e-mail. While a detailed study of participant's e-mail and file organisation habits is outside the scope of this evaluation, this finding suggests the need to support people interacting with personal information stored in different organisation schemes across multiple tools (even though the information may be related).

DISCUSSION

The evaluation shows that people can effectively access and organise a small set of personal information with our tabletop interface. While the organisation of information was done by placing objects in Storage Bins, these bins could be used to assign tags to objects, or to place them in a folder. We saw that the focus item approach to personal information navigation was well received, and was intuitive to use on a tabletop—it uses a natural dwell gesture, and does not require keyboard input for attribute and content-based information access. Given that a collaborative tabletop PIM interface will need to support interaction with multiple people's information concurrently, it must support access that is not dependant on a particular storage technique (e.g. hierarchical structure or tagging) or a personal organisation scheme. The navigation techniques exhibited in the Focus interface are promising for supporting collaborative PIM with a multi-user tabletop interface.

The evaluation highlighted the issue of people's ability to identify personal informational when it is presented in a purely graphical form on a tabletop. We saw that participants immediately drew their attention to information that contained images (such as photos, PDF documents and webpages). As e-mails tend to look highly similar on the tabletop when they presented at a small size, this finding highlights the need for designers of tabletop PIM interfaces to explore ways to distinguish between representations of textual information. While a filename or e-mail subject could be presented to help users distinguish between e-mails that are similar in appearance, this is challenging because of the properties of the tabletop medium—the text would need to be sufficiently large to read on the tabletop (even when the e-mail is presented in a small thumbnail view), but this could potentially increase the degree of clutter in the interface (an

important and challenging problem to address in tabletop interface design). Familiarity with the personal information will also have an impact on a person's ability to identify it on the tabletop, though the issue of supporting people interacting with unfamiliar information (as was the case in our evaluation) is important to address in collaborative PIM—not only do we want to support people interacting with their own personal information, we also want to support interaction with other people's personal information using a shared workspace.

As the Focus interface presents information on the tabletop without conveying its storage structure, there is the potential for the interface to become cluttered when too many objects are displayed on the tabletop at once (e.g. when there are many items related to a selected focus item). The results suggest that clutter was sufficiently managed by our tabletop interface, using the focus object technique (only showing items relevant to a selected focus item), the Black Hole for temporary removal of objects, and sizing information according to its relevance to the focus item. No participant mentioned clutter as an issue when completing the evaluation tasks. However, the evaluation highlighted the problem of overlapping objects—particularly when people created groupings of objects that conflicted with groupings made for other focus items. This is an important issue to address in future work.

Our questionnaire data indicated that, from our small sample size, people did not use the same organisation scheme for both their file system and e-mail collection. While this finding does not indicate that most people organise their file system and e-mail differently, it indicates that if collaborative PIM is going to be used by a range of people, for a range of purposes, we must support tabletop interaction with information that is organised differently (both at a storage level and a user interface level) and managed with different tools.

A critical issue in collaborative PIM, raised by all participants in their questionnaire responses, is privacy. A key factor to the uptake of collaborative PIM interfaces is how people are able to control privacy, so that they have complete control over the information available at the tabletop, and who can access it. While the Focus Exporter was designed with privacy in mind (it allows strict control over which tabletops can connect to the computer, and the information accessible to them), the Focus Exporter was not shown or discussed with participants as it was not a factor in the evaluation. This finding indicates that people are generally not going to access all of their personal information at a tabletop in a collaborative setting, and designers of future collaborative PIM systems should take into account that people are likely only to export a small subset of their information (possibly related to a single project or issue that they must discuss at the tabletop).

FUTURE WORK

There are many interesting enhancements that we wish to make to the Focus interface. An important issue for Focus (and tabletop interfaces in general) is clutter management, and we will explore more techniques to reduce clutter and support users managing it. We will adopt a thread-based presentation of e-mails on the tabletop in order to reduce the amount of space taken by similar e-mails. A possible approach—that is consistent with our existing interface—is to have past e-mails of a thread attached to the back of the most recent e-mail. This means that a user flips the e-mail over to see the related e-mails. Our evaluation highlighted the issue of overlapping objects in Focus. We will explore techniques for allowing users to easily align a group of objects so that they do no overlap.

Focus currently supports the e-mail start view and the hierarchical start view (where the first file in each directory, and the most recent e-mail in each mailbox is displayed). We will explore other techniques for providing an overview of the personal information available at the tabletop. For example, providing a list of tags in a faceted browsing interface may be an appropriate way to allow users to browse their information in order to make an initial focus selection.

The Focus interface currently supports e-mails, stored webpages, PDF documents and photos. There are other types of personal information that would be useful to support in a collaborative PIM system. In particular, calendar events and address book contacts would provide useful focus selections—a user would be able to easily see all information related to a particular event, or all information related to a selected person. When presenting multi-page documents in the Focus interface, currently only the first page is displayed. In future, we will provide support for navigating through multi-page documents. For example, the document could be selected from its corner and turned over to view the next page (to be consistent with the physical nature of the tabletop interface).

Finally, a key area for future work is conducting evaluations involving multiple people collaborating with their own personal information. This will involve people conducting real tasks, with personal information stored on their own computers.

CONCLUSION

We have presented Focus, a collaborative PIM interface that supports multi-user tabletop interaction. This interface embodies a novel exploration of collaborative PIM using shared, multi-user devices. We have conducted a preliminary evaluation, which indicated that Focus shows promise for supporting navigation of remote collections of personal information, within the constraints of collaborative tabletop interaction. The evaluation highlighted the need to enhance support for clutter management, and the need for display techniques that help people distinguish between textual documents that look highly similar (such as e-mails). Reactions to the concept of collaborative PIM at a tabletop were positive, though privacy control is an important issue that will affect whether people are willing to use collaborative PIM interfaces in the real world.

ACKNOWLEDGEMENTS

We thank our experiment participants for taking part in our evaluation.

REFERENCES

1. Agarawala A. and Balakrishnan R. Keepin' it real: Pushing the desktop metaphor with physics, piles and the pen. In *Proceedings of CHI '06*, pp. 1283–1292. ACM Press, New York, NY, USA (2006).

2. Apted T., Kay J., and Assad M. Sharing digital media on collaborative tables and displays. In *Online Proceedings of The Spaces In-between: Seamful vs. Seamless Interactions (in conjunction with UbiComp 2005)* (2005).

3. Apted T., Kay J., and Quigley A. Tabletop sharing of digital photographs for the elderly. In *Proceedings of CHI '06*, pp. 781–790. ACM Press (2006).

4. Boardman R. and Sasse M.A. "Stuff goes into the computer and doesn't come out" A cross-tool study of personal information management. In *Proceedings of CHI '04*, pp. 583–590. ACM Press, New York, NY, USA (2004).

5. Boardman R., Spence R., and Sasse M.A. Too many hierarchies? The daily struggle for control of the workspace. In *Proceedings of HCI International 2003*, pp. 616–620. Lawrence Erlbaum Associates, New Jersey, USA (2003).

6. Collins A., Apted T., and Kay J. Tabletop file system access: Associative and hierarchical approaches. In *Proceedings of TABLETOP '07*, pp. 113–120. IEEE Computer Society, Washington, DC, USA (2007).

7. Cutrell E., Robbins D., Dumais S., and Sarin R. Fast, flexible filtering with Phlat — personal search and organization made easy. In *Proceedings of CHI '06*, pp. 261–270. ACM Press, New York, NY, USA (2006).

8. Dumais S., Cutrell E., Cadiz J., Jancke G., Sarin R., and Robbins D.C. Stuff I've Seen: A system for personal information retrieval and re-use. In *Proceedings of SIGIR '03*, pp. 72–79. ACM Press, New York, NY, USA (2003).

9. Dumais S., Cutrell E., Sarin R., and Horvitz E. Implicit queries (IQ) for contextualized search. In *Proceedings of SIGIR '04*, pp. 594–594. ACM Press, New York, NY, USA (2004).

10. Erickson T. From PIM to GIM: personal information management in group contexts. In *Commun. ACM*, 49(1):74–75 (2006).

11. Everitt K., Shen C., Ryall K., and Forlines C. MultiSpace: Enabling electronic document micro-mobility in table-centric, multi-device environments. In *Proceedings of TABLETOP '06*, pp. 27–34. IEEE Computer Society, Washington, DC, USA (2006).

12. Freeman E. and Gelernter D. Lifestreams: A storage model for personal data. In *SIGMOD Record*, 25(1):80–86 (1996).

13. Gemmell J., Bell G., Lueder R., Drucker S., and Wong C. MyLifeBits: Fulfilling the Memex vision. In *Proceedings of MULTIMEDIA '02*, pp. 235–238. ACM Press, New York, NY, USA (2002).

14. Hartmann B., Morris M.R., and Cassanego A. Reducing clutter on tabletop groupware systems with tangible drawers. In *Adjunct Proceedings of UbiComp 2006* (2006).

15. Karger D.R., Bakshi K., Huynh D., Quan D., and Sinha V. Haystack: A customizable general-purpose information management tool for end users of semistructured data. In *Proceedings of CIDR 2005*, pp. 13–27 (2005).

16. Karger D.R. and Jones W. Data unification in personal information management. In *Commun. ACM*, 49(1):77–82 (2006).

17. Morris M.R., Paepcke A., and Winograd T. TeamSearch: Comparing techniques for co-present collaborative search of digital media. In *Proceedings of TABLETOP '06*, pp. 97–104. IEEE Computer Society, Washington, DC, USA (2006).

18. Robertson G., Czerwinski M., Larson K., Robbins D.C., Thiel D., and van Dantzich M. Data Mountain: Using spatial memory for document management. In *Proceedings of UIST '98*, pp. 153–162. ACM Press (1998).

19. Ryall K., Morris M.R., Everitt K., Forlines C., and Shen C. Experiences with and observations of direct-touch tabletops. In *Proceedings of TABLETOP '06*, pp. 89–96. IEEE Computer Society, Washington, DC, USA (2006).

20. Scott S.D., Carpendale M.S.T., and Habelski S. Storage bins: Mobile storage for collaborative tabletop displays. In *IEEE Computer Graphics and Applications*, 25(4):58–65 (2005).

21. Shen C., Everitt K., and Ryall K. UbiTable: Impromptu face-to-face collaboration on horizontal interactive surfaces. In *Proceedings of UbiComp 2003*, vol. 2864 of *Lecture Notes in Computer Science*, pp. 281–288. Springer, Berlin/Heidelberg (2003).

22. Shen C., Lesh N.B., Vernier F., Forlines C., and Frost J. Sharing and building digital group histories. In *Proceedings of CSCW '02*, pp. 324–333. ACM Press, New York, NY, USA (2002).

23. Vernier F., Lesh N., and Shen C. Visualization techniques for circular tabletop interfaces. Tech. Rep. TR2002-001, Mitsubishi Electric Research Laboratories (2002).