

A 'Pile' Metaphor for Supporting Casual Organization of Information

Richard Mander, Gitta Salomon and Yin Yin Wong

Human Interface Group, Advanced Technology
Apple Computer, Inc.
20525 Mariani Ave., MS 76-3H
Cupertino, California 95014
(408)996-1010

ABSTRACT

A user study was conducted to investigate how people deal with the flow of information in their workspaces. Subjects reported that, in an attempt to quickly and informally manage their information, they created piles of documents. Piles were seen as complementary to the folder filing system, which was used for more formal archiving. A new desktop interface element – the pile – was developed and prototyped through an iterative process. The design includes direct manipulation techniques and support for browsing, and goes beyond physical world functionality by providing system assistance for automatic pile construction and reorganization. Preliminary user tests indicate the design is promising and raise issues that will be addressed in future work.

KEYWORDS: interface design, design process, interactive systems, user observation, desktop metaphor, interface metaphors, pile metaphor, information visualization, information organization, end-user programming.

INTRODUCTION

As the amount of information users confront on their computers increases, tools to organize and manipulate this information become increasingly important.

Today's direct manipulation computer interfaces, such as the Macintosh® desktop interface [1], offer limited means of handling information. Users can manually place files within folders, organized in a rigid hierarchy. Users are responsible for appropriately filing all items; the system offers little assistance in this often tedious task. Recent enhancements, such as "aliases" [2], allow users to overcome a frequent problem, namely that an item belongs in more than one folder. However, the folder as the sole container type presents an impoverished set of possibilities.

The real world provides a rich array of organization systems. In the past, researchers have looked at how users find items in their physical offices [9]. We conducted a study to observe how users *organize* the large amounts of information they work with in their physical offices. Our study differed from previous work in that we looked at ways in which people use and interact with filing systems.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

We were also interested in how people work with assistants when dealing with information.

By examining individuals' information management schemes, we were able to extract and extrapolate a number of interesting interface ideas for a graphical interface. Our intent was not simply to emulate physical world functionality – several investigators have argued against this procedure [3,6] – but rather to leverage users' knowledge to create an intuitive and powerful system that goes beyond physical world capabilities. Using this approach, we sought to construct a design which provides new functionality and enhances the user interface.

Like Malone [9], we found that users like to group items spatially and often prefer to deal with information by creating physical piles of paper, rather than immediately categorizing it into specific folders. Computer users are confronted with large amounts of information, but currently are only provided with a hierarchical filing system for managing it.

Therefore, we propose that incorporating 'piles' within a graphical user interface could provide a number of interesting possibilities. Users have difficulty deciding where to file a new item; piling requires less mental effort. Today's office assistants use piles as a way of suggesting categories to others; computerized agents might make use of them in the same way to convey a certain degree of imprecision in the suggested organization. Piles may also provide an appropriate representation for the results of information retrieval algorithms which are inherently inexact [13].

At least one system, BUSINESS [11], previously explored this interface metaphor as a construct within a text-based application programming language. For example, a user could initiate an action by typing an instruction such as "Empty the In Box onto the Work Pile." However, there was no graphical representation, and so the system could do little more than allow the user to issue programmatic commands using a subset of English.

This paper provides both specific design ideas and insight into our design process as it progressed from user interviews to design to testing. The first section describes findings from observing and interviewing office workers. In particular, we report why folders were not always appropriate and how and when users found piles useful. We then describe the interface designs inspired by these observations. In the third section, we report results from informal tests of

these designs. In conclusion, we describe directions for future work.

USER INTERVIEWS

As part of our design process, we undertook a user study to find out how people deal with information in their physical workspaces. Studies of this kind are important in helping us understand the user's perspective. Our aim was to identify aspects of the real world work process which could offer insight into a new, more powerful interface.

Method

The study involved interviews with thirteen men and women in Marketing, Support, Human Resources, and Technical departments within Apple Computer, Inc. The interviews lasted between 30 and 60 minutes and were conducted in the participant's work area. All interviews were videotaped.

We asked people to describe the way information arrived in their work area, what they initially did with this information, where it went next, and how it was finally stored. Participants gave us a tour of their workspace to help us understand what purpose various cabinets, shelves, and storage devices served. We also took a small pile of documents with us and asked participants to judge what these documents were on the basis of their appearance. In this way, we could find out how they worked with completely unfamiliar information. Since we are interested in developing ways for the computer to help the user, we also asked people how they worked with assistants.

Findings

Our subjects used a variety of techniques – folders, file cabinets, file racks, piles, binders, card files, and bulletin boards – for managing the information in their offices. Since our primary concern in this paper is the uses of folders and piles, we'll focus on observations relevant to these items.

Uses for folders. File folders were used in several ways. As could be expected, items were placed in folders which were in turn placed in file cabinets as a means of archiving information not currently needed. Users applied a variety of organizations to their file cabinets, ranging from totally random arrangements to strict alphabetical and color coded systems.

Users were sometimes dissatisfied with using folders in this way, because they were required to make an explicit decision about how to categorize individual items. This was often especially difficult with new information. One user said "I'm not always as good at categorizing things as I would like...it's hard to get it right and I'm sort of a perfectionist, so I think that I should know exactly how I should do it...I like things in their place, but I can't figure out exactly what place."

One solution, identified by several users, would be to file the information in several places. However, even though copiers were near at hand, people did not choose to duplicate information in order to store it in more than one folder.

Folders were also used in informal ways. Many people mounted folders in racks, which enabled the folders to stand up. These folders were used for frequently accessed information – most often action items and items requiring

regular maintenance, such as expense reports and things to read. Some users ordered or changed the orientation of the folders in their racks to make the most important or urgent information prominent. Folders were also used as a storage medium within piles, as a way to hold together a certain group of items. As one user commented, "...[I] folderize to keep things neat...there's no hierarchy in there, because building a hierarchy takes too much time."

Based on these observations, we inferred two things about the current folder-based interface offered on the Macintosh: the categorization problems are presumably amplified by the use of multiply nested folders, and support for more informal grouping techniques, such as racks, could be useful.

Piles: A less rigid categorization system. In addition to using folders, users grouped items into piles. For example, most workers kept information they needed in a specific working area. A common strategy was to create separate piles for each project and place them within the working area, at distances that reflected their urgency. Many workers also created piles for incoming information that they could not deal with immediately. The contents of users' piles was clearly not restricted to paper documents – we observed piles composed of various items such as books, folders, reports, binders, cassette tapes, video tapes, postcards, envelopes, magazines, journals, and boxes.

People used piles instead of hierarchical folders because they did not require detailed categorization and they could be more easily reordered than a folder and file system. For many workers, the pile was viewed as an entity that was subject to change. Users reported that over a period of time, items within a pile would often be reshuffled and broken down into several sub-piles, and an informal process of categorization would begin. We noted several approaches to separating material within piles: some users stacked materials at different angles, while some placed dividers within the pile.

To the outside observer, an office containing piles often appears disorganized. However, all of our participants had several piles in their workspace and in most cases, they knew what was in each pile and could tell us quite a lot about its history. Seemingly disordered piles were often sensible to the person who created them, because they developed through many interactions over a long period of time. For instance, many piles grew as newer items were added to the top, and workers could tell where things were by their date, since the stack was ordered chronologically.

Piles: self-revealing, browseable. Several users remarked that the outer appearance of their piles conveniently allowed them to recognize particular items. Our subjects were also able to make use of the appearance of previously unseen piles. We asked them to look at a small pile of unfamiliar materials which we took with us to the interview. By looking at the pile's outside form, they were able to infer quite a lot about its contents.

Consequently, we noted that piles facilitate browsing, and we observed four different browsing methods. In the *edge* browse method described above, people looked at the outside edges of the pile for clues about the items within. Information such as color, texture, and thickness was commonly used to judge the contents of a pile. In the *restack*

method, people started at the top of the pile and dealt with each item in turn by lifting it off the pile, looking at it and then placing it somewhere other than back on the pile. In the *hinge* method, the items stayed in the pile, but the pile was hinged open at different points to display a single item. The final method was to *spread out* a pile and look at its contents in parallel.

Assistance with information management. Most participants did not have an assistant, but said they would welcome one. We asked those who did have assistants to describe how they worked together.

Assistants commonly took care of routine tasks, such as sorting mail into different categories. For people who had to deal with large amounts of information, their assistant acted as a filter, passing along urgent material and removing junk mail. Some assistants would reorganize the workspace and create a filing system in which information could be more easily organized. This usually happened in collaboration with the worker. Typically the assistant would suggest categories and discuss these with the worker before actually filing the material. The assistant would often not understand the technical content, but could scan through the materials looking for keywords that might help in the categorization task. Piles were often used by assistants to indicate potential categories. As one assistant remarked, "I'll go into his office and put [labels] on piles on his floor and he'll look at it and say 'no' or he'll say 'that's pretty good'."

FROM OBSERVATION TO DESIGN SKETCH

The next step in our process was to take our observations and develop a number of design sketches using MacroMind's Director™ application [8] which supports scripted interaction and animation. These design sketches illustrated particular interaction techniques and were used to facilitate group discussion about interaction possibilities and the technology necessary to support them. They centered around the development of a new organizational element – the pile – which would support informal groupings of items on the computer desktop. In addition, we extended the metaphor to include functionality which could only be provided by the existence of a computer. The design sketches created are described below.

User-created piles. One objective was to allow users to create piles of mixed content and multiple data types. Each item within a pile would be represented by a miniature depicting its first page and extent (see Figure 1a). We wanted to maintain the informal quality of physical piles by provid-




Figure 1. **Piles on the desktop.** In general, piles can contain various media, such as folders and individual documents. The pile in (a) was created by the user, and is consequently disheveled in appearance. In addition, the system can create piles for the user, based on rules explicitly stated by the user or developed through user-system collaboration. These piles have a neat appearance, as shown in (b), to indicate that there is a script, or set of rules, behind them.

ing direct manipulation techniques which resemble real world interactions. For example, a pile is created by overlapping two items; items are added to an existing pile by simply placing them on top (Figure 2). These user-created piles have a disheveled look.

System-created piles. In addition, we postulated that the system could create piles for a user. As shown in Figure 1 (b), these piles would have an orderly appearance. The system would assemble these piles using a script either developed through user-system collaboration, or explicitly written by the user. By creating and maintaining piles for the user, the system could serve as an office assistant.

How might this user-system collaboration work? Potentially, the user could supply sample documents as input for pile construction. By analyzing these documents, the system could offer various criteria for script construction. For example, the system could determine a document's unique terms and let the user select the specific terms to use as piling criteria. Additionally, the system might extract structural data, such as the "Re:" line in a mail message and ask the user if similar mail messages should be collected into the pile. Malone suggested a similar tact for automatic classification [9] and reported successful results in the Information Lens system [10]. As shown in Figure 3, our design provides a way for users to gradually learn to create scripts. As in the work of MacLean et al [7], we wanted to provide a natural way for users to approach "tailorability" of piles as a part of the system.

Support for browsing. We wanted to support some of the browsing techniques users applied in their physical offices.




Figure 2. **Adding a document to a pile.** If a document is positioned over an existing pile, the pile highlights to show that it can accept the new document. When the mouse button is released the document 'drops' onto the pile.




Figure 3. Scripting a pile. (a) depicts a mail area containing two scripted piles; one for important items, one for everything else. Over time, the criteria for 'important' may change. As shown, an item in the 'other' pile has been removed because the user desires that it, and items like it, now appear in the important pile. When this item is dropped onto the important pile, as shown in (b), the system queries the user to find out whether this action is a singular event or whether the pile's script should be modified. If the user chooses to modify the script, the system suggests criteria which could be used, as shown in (c). Alternatively, users can gain direct access to the scripting language and write their own criteria via the "Script..." button. Once the script is updated, items satisfying the new criteria visibly move to the 'important' pile.




Figure 4. Browsing by spreading out a pile. Gesturing sideways with the mouse pointer, or with a finger in the case of a touch screen, causes the pile contents to spread out. Individual items can now be directly manipulated.




Figure 5. Browsing while maintaining the pile's structure. Gesturing vertically with the mouse pointer as shown in (a), or with a finger in the case of a touch screen, generates a 'viewing cone' (b) that contains a miniature version of the first page of the item under the pointer. This viewing cone will follow the vertical position of the pointer; the miniature changes as the pointer moves over each item. The user can move through the pages of an item in the viewing cone by using the left and right cursor keys on the keyboard. When an item is visible in the viewing cone, it can be selected by clicking the mouse button. The item then appears next to the pile on the desktop, as shown in (c).




Figure 6. *Visualizing a pile's contents.* The pile shown is within a 'visualizing environment' that allows the user to select and visualize several criteria. Criteria can be mapped to the pile's order, the color of the items within the pile, or the way a pile is broken into sub-piles. In (a) the pile is both ordered and colored by date. In (b) the user chose to 'pile by' content. Therefore, the system separated the original pile into four content-based piles. Three are labeled with specific terms suggested by the system (e.g. "architecture"), appear neat and are now scripted to maintain similar content. The remaining disheveled pile, "other," contains items which did not fit into any of the other three piles.

By virtue of using miniatures of the actual documents, we offered edge browsing capabilities. In addition, we explored gestural inputs as a way to invoke other browsing methods. For example, a horizontal gesture would spread out a pile so that miniatures of each item's first page were visible (Figure 4). A vertical up-and-down movement over a pile would allow users to browse a pile using a 'viewing cone' (Figure 5). When an item was visible in the viewing cone, the user could move through miniature representations of its pages by using the cursor keys on the keyboard. In the design sketches, a mouse was used to create the gestures, but we thought that these interaction techniques would be particularly well-suited to a touch screen display.

Managing piles. In physical offices, the user is confronted with many pile management tasks, such as re-piling and sub-piling when a particular pile becomes unwieldy or specific information must be retrieved. We wanted the system to act as a collaborator in dealing with these issues, and therefore designed a 'visualizing environment' which would help users understand the contents of piles. As shown in Figure 6, a user might choose to emphasize certain criteria in a pile by using order, color or sub-piles. A user can elect to view combinations of criteria simultaneously. For example, the user could choose that items in a pile be ordered by date. The user might also request that the items in a pile be color coded according to their data type. Additionally, the user might have the system suggest subject-based sub-piles, by choosing the "pile by content" option. The sub-piles deemed useful could be moved out of the visualization area for use on the desktop.

TESTING USER'S EXPECTATIONS OF PILES

The design sketches raised interest amongst our colleagues and were the focal point for discussions. However, since the sketches were 'hard-wired,' we did not know if the interaction techniques were usable and of value to end-users. Consequently, we undertook a user test of the interaction techniques.

We constructed a suite of prototypes in Director that supported the interactivity we wished to test. We hoped to gauge people's expectations about the inclusion of piles on the desktop. Our method was informal, resembling the type of testing described in [4,12], in order to provide us with quick results that could be used in design iteration.

Method

Five men and five women in nontechnical positions at Apple Computer were individually tested in approximately one hour sessions. The subjects were asked to think aloud [5] while working through 5 tasks, and the sessions were videotaped. The first two tasks compared two different pile models. In the third task, users explored methods of initiating pile browsing. Task four allowed users to indicate preferences between three different viewing cone representations. In the final task, users were asked to locate items within a pile. At the conclusion of the test, users were informally asked for their comments and general impressions concerning piles.

Piling models. Two different models for a pile were compared: a "document-centered" model and a "pile-centered" model. Possible ordering effects were avoided by varying the presentation of these two models across users.

In the "document-centered" task, the pile was represented as a collection of individual items. The user was presented with a series of colored rectangles within a white screen area. These rectangles were intended to represent files on a desktop. The rectangles could be selected and moved with the mouse. When one rectangle was placed over another rectangle, both would fall back to create a disheveled pile. Additional documents could be added to an existing pile by moving them over the pile and releasing the mouse. Documents could be removed by individually selecting them via any visible region and dragging them away from the pile. The pile itself could not be moved as a unit.

In the "pile-centered" task, piles were created in the same way, except that the pile acted more like a Macintosh folder – a single entity containing a collection of documents. When one document 'rectangle' was moved over another rectangle, the latter would highlight to indicate a pile would be formed if the mouse button was released. Subsequent documents moved to the pile would automatically drop onto the top of the pile. The pile itself, as opposed to independent documents, could then be dragged around the desktop by mouse-clicking on any part of it.

Initiating browsing. In this task, participants tried out different ways of initiating pile browsing. They compared double-clicking and the horizontal gesture (shown in Figure




Figure 7. Test of different viewing cone representations. Users were presented with three different ways a pile could react during viewing cone browsing. In each case, the viewing cone contains a miniature of the first page of the document being examined. In style (a), the pile remains stationary. In (b), each item above the item being currently viewed is moved to the left side. In (c), the item currently being viewed temporarily moves out of the pile and to the right. Users preferred methods (b) and (c).

4) as ways to spread out a pile's contents. They also compared double-clicking and vertical gesturing (Figure 5) for initiating browsing with the view cone.

Viewing cone representations. In this task, participants were presented with three different visual representations of the viewing cone (see Figure 7). All were initiated by the same interaction – a single mouse click on the pile – but the order of presentation was varied for each user. When users settled on a preferable representation, they were shown how to use the keyboard to examine a miniature document's pages while it was within the viewing cone. Users were shown which key would move the document forward one page, and which would move it backwards.

Finding items within a pile. In this task, participants were asked to use the viewing cone and paging ability to locate specific pages within documents in the pile. First they were asked to locate a picture of a hand on a mouse, for which they were shown a real report illustration as a stimulus. Then they were asked to locate three separate items within the pile: a colored bar graph, a map of North America which was contained in an Atlas document, and a document containing bullet point text. Users were not timed – this part of the test was aimed at determining if pile browsing was, in general, qualitatively suitable to users for locating information.

Results

Piling models. Although each user had a clear preference for one of our methods of pile creation (“pile-centered” or “document-centered”), neither method was judged to be clearly superior. In the “document-centered” model, users liked the ability to grab an individual document within a pile. A problem with this model was that users were not sure how to move a pile as a unit, since selecting any part of the pile led to moving an individual item rather than the pile as a whole. In the “pile-centered” model, users liked the way the system automatically aligned the items in the pile, the ability to move a pile as a unit, and the highlighting that indicated a pile was ready to accept an item. A problem with this model was the difficulty of selecting an individual item within the pile.

Most users also expected that any desktop item could be added to a pile. This led to discussion of what would happen if a document was placed on top of an isolated folder; users were unsure whether the item would go into the folder or if a new pile would be created. Most users thought

that, based on their previous Macintosh experience, the item would go into the folder. This raises questions about how the pile metaphor fits into the current Macintosh desktop metaphor.

Most users asked for features generally available in desktop systems, but which were not present in the testing prototypes. For example, they wanted to be able to add a selected group of items to a pile, name piles, apply ordering schemes based on date, size, name, and kind, and control where a document was placed within a pile.

Since users liked and disliked certain features of each model, new design work will be undertaken to create models that both embody users' preferences and are internally consistent. Further testing of these new models will be conducted.

Initiating browsing. Subjects tried using both gestures and mouse double-clicks to spread out a pile and also to obtain the viewing cone. In both cases, 9 out of the 10 participants preferred the double-click method. They found it faster and more Macintosh-like, which was not unexpected given that the subjects were all accustomed to the Macintosh. However, users also felt that the gestures were non-intuitive and somewhat ambiguous, and that the piles might be spread out accidentally while moving the cursor around on the screen. The gestures were originally intended for use on a touch screen and most participants said that using a finger on the screen for the gesture might be more intuitive than using a mouse. This needs to be confirmed in a test with a touch screen. Note that we did not ask users *which* of the browsing methods they would want initiated by the double-click action; we only ascertained that they preferred double-clicking over gesturing.

In general, users thought they would make use of the ‘spread out’ view. Since all items were visible at once, it supported recognition and comparison. A few users expressed interest in viewing a grid layout rather than the overlapping one used in our testing prototype. While in this view, most users expected to both be able to act on individual items in standard ways, and move the documents as a group. In addition to the miniature representation of each item, many users requested that other information such as name, date, and kind be made available, and that the system provide representations which would specifically help the user differentiate similar items.

Viewing cone representations. Of the three viewing cone designs, the stationary pile version (Figure 7a) was rejected by all 10 users. All thought it was difficult to gauge where they were within the pile. Four users preferred the 'side shift' style (Figure 7b), 5 preferred the 'document pull-out' style (Figure 7c), and 1 user was undecided between these latter two designs. Both of the preferred designs clearly provided a view of an item's location in the pile, in addition to a representation shown within the viewing cone. Although it was not implemented in the prototype, once users had an item visible in the cone they often tried to grab it by either releasing and then quickly clicking the mouse, or by attempting to drag it from the pile.

Most users liked the viewing cone as a browsing method. It made it possible for them to identify items by their miniature representation without disturbing the pile's state. Users also liked the ability to view any page of an individual item, although not all were pleased with using the cursor keys on the keyboard to cause this action. Page numbering information (e.g. '1 of 10') was found valuable while paging through the document, because it indicated the relative size of each item, as well as position within an item. One user desired random access to any page via selection of its number from the keyboard. A few users expressed interest in being able to target the viewing cone at any item on the desktop – a single document, a folder – and not just items in piles.

During the tests we noted a potential problem with the viewing cone implementation – users might need to depress the mouse button for a long time while browsing, which could lead to repetitive stress injury. A possible solution is to invoke the cone whenever the user clicks on a pile, thereby alleviating the need for the mouse button to be continuously depressed. This would also allow the user to click the mouse button again to select an item for removal from the pile while the viewing cone was active.

Finding items within the pile. We showed the subjects a physical version of a report and identified a specific illustration which we wanted them to find within a pile on the computer desktop. Users were asked to use the viewing cone and cursor keys to examine items in the pile. All of the users were able to find the illustration within a reasonable amount of time. As mentioned earlier, we were not concerned with timing information, but rather with the feasibility of the viewing cone for this task.

Since the picture was within a report which was bound in a green-edged cover, several users recognized the document within the pile by its clearly visible green spine. A common strategy was to subsequently move through the document's miniature pages, looking for a small colored image in the top right corner of a page. Only one user took advantage of the page numbers on the miniature representations to identify the page. A few users did not expect the document on the computer to have the green spine that was present on the physical report because they perceived it to be an addition which the system could not have known about. These users' strategy was to start at the top of the pile and systematically look through every item, page by page, to find the picture.

We also asked the subjects to find a colored bar graph, a map of the North American continent within an atlas, and

some bullet point text. Only some of the users found the items, and with some difficulty. Many users felt they would do better with their own information and their difficulty was due to a lack of familiarity with the material in the pile. Several users discussed ways they would like the system to help them in such a situation. They commonly wanted the ability to search for specific data types, names, keywords, and other identifiers.

From this feedback, we inferred that it might be useful to give the user control over the information presented in the viewing cone. For instance, when searching for a graph, the user could select data type 'graph' as the search criteria, thereby causing the viewing cone to display only pages containing graph data types. This could be a powerful way to search, since it would enable the user to tailor the view according to current needs.

General discussion. At the end of the test, we asked users how they would use piles, and how the system might assist them. In general, users were receptive to the idea of having the system help them with their routine tasks, such as sorting incoming mail. Most users reported having between two and five mail systems, fax, and voice mail. They liked the idea of receiving all incoming information in a pile which could be accessed with the viewing cone. Within such a pile, they would want the system to prioritize items using characteristics such as sender, topic, content keywords, date, and urgency. We anticipate these priorities could be learned by the system over a period of time by watching the user interact with incoming information.

FUTURE WORK: FROM DESIGN SKETCH TO IMPLEMENTATION

There are many areas in which this work can proceed. A few of our current directions are described below.

Improving Designs and Working with Familiar Data

We plan to further explore the appropriate model for a pile – and how to possibly combine users' expectations about its document-centeredness vs. pile-centeredness – by iterating on our previous Director prototypes.

In addition, we intend to build prototypes that incorporate items of relevance to the individual being tested. The informal tests described above involved fabricated data that was unfamiliar to our subjects. In order to continue refining our designs, we need to construct prototypes that will allow subjects to interactively use piles for their own information over an extended period of time. An extension to the current Finder interface that would allow users to create and work with piles alongside folders would provide an excellent opportunity to further these designs. However, it may prove more feasible to undertake the next round of iteration by addressing a limited domain, such as a mail system.

Browsing by Other Criteria

The current design allows users to browse the contents of piles by viewing miniature representations of each item. While users found this feature useful in the tests, they also expressed interest in accessing other representations. We are currently exploring the types of "browse by..." criteria the system might offer. For example, users might want to selectively emphasize some data type during browsing, as in the case of 'show me all the documents containing movies within this pile.' When confronted with unfamiliar

data, users might want to browse by textual abstract, since a miniature visual representation might not provide insight into an unknown item's content.

Technology to Support Pile Interactions

The interface designs described in this paper were primarily inspired by observations with users, and not necessarily by existing technology. At the time of design, we were unsure if information retrieval techniques could adequately support some of the interactions, such as pile scripting-by-example or sub-pile creation. Consequently, we initiated a collaborative research effort with the Information Retrieval Team within Apple Computer's Advanced Technology Group.

Some preliminary work in implementing low level support for pile functionality has been undertaken. Current research is focussing on a clustering technique that would automatically create sub-piles. For example, a user could supply the system with a pile of documents, and based on the content of the documents within that pile, the system would suggest and describe suitable sub-piles.

As this work progresses, we plan to adapt our designs to reflect the technology that can be realized.

ACKNOWLEDGEMENTS

We would like to thank Dan Rose and Tim Oren for exploring information retrieval systems that will support sub-piling and other pile management operations; Stephanie Houde for creating the Director prototypes used in testing; Penny Bauersfeld and Leo Degen for their participation in early design sessions; and Tom Erickson for input on the user study design and feedback on the various drafts of this paper.

BIBLIOGRAPHY

- [1] Apple Computer, Inc. *Human Interface Guidelines: The Apple Desktop Interface*. Addison-Wesley Publishing Company, Inc., Reading, MA, 1987.
- [2] Apple Computer, Inc. *Inside Macintosh, Volume VI*. Addison-Wesley Publishing Company, Inc., Reading, MA, 1991.
- [3] Cole, I. Human aspects of office filing: Implications for the electronic office *Proceedings of the Human Factors Society, 26th Annual Meeting*, Seattle, Washington. 1982.
- [4] Gomoll, K. Some Techniques for Observing Users *The Art of Human-Computer Interface Design* (ed. Brenda Laurel) Addison-Wesley Publishing Company, Inc., Reading, MA. 1990. pp. 85-90.
- [5] Ericsson, K.A. and Simon, H. A. *Protocol analysis*. Cambridge, Massachusetts: MIT Press. 1984.
- [6] Lansdale, M. The psychology of personal information management. *Applied Ergonomics*, 55, (1988), pp. 55-66.
- [7] MacLean, A., Carter, K., Lovstrand, L. and Moran, T. User-Tailorable Systems: Pressing the Issues with Buttons. In *Proceedings of CHI 1990 (Seattle, Washington, April 1-5, 1990)* ACM, New York, 1990. pp. 175-182.
- [8] MacroMind, Inc. *Director™ 2.0*. April 1990.
- [9] Malone, T. W. How do People Organize Their Desks? Implications for the Design of Office Information Systems. *ACM Transactions on Office Information Systems*, 1,1, (January 1983), pp. 99-112.
- [10] Malone, T. W., Grant, K.R., Turbak, F.A., Brobst, S.A. and Cohen, M.D. Intelligent Information-Sharing Systems. *Communications of the ACM*, 30, 5, (May 1987), pp. 390-402.
- [11] Miller, P., Tetelbaum, S. and Webb, K. BUSINESS – an end-user oriented application development language. *SIGMOD Record*, 12, 1, (October 1981), pp. 38-69
- [12] Nielsen, J. Usability Engineering at a Discount. In G. Salvendy and M.J. Smith (Eds.), *Designing and Using Human-Computer Interfaces and Knowledge Based Systems*. Amsterdam: Elsevier. 1989.
- [13] van Rijsbergen, C.J. *Information Retrieval. (2nd. Ed.)* Butterworths, London, England. 1983.