

# Desktop History: Time-based Interaction Summaries to Restore Context and Improve Data Access

Sean Uberoi Kelly and John P. Davis

Social Computing Group, Microsoft Research  
One Microsoft Way, Redmond, WA 98052-6399 USA  
{seankell, jpdavis}@microsoft.com  
+1 425 882 8080

## ABSTRACT

The Desktop History Tool presents users with a summary of the data they worked with throughout past days and weeks. The user's own interaction history provides an invaluable self-organizing structure for the desktop interface. A user study was conducted to assess whether users could better discover and reconstruct important contexts by browsing and searching with desktop history, compared to the methods they typically used. We found that the history tool enabled users to discover a greater amount of important data and discover related data, allowing them to develop more thorough recreations of previous workspaces. We conclude that a time-based presentation of a user's interaction history is a useful tool for supporting richer context-aware interfaces.

## Keywords

Calendar, desktop, time-based computing, chronology, summarization, context, interaction history

## INTRODUCTION

Current desktop user interfaces can hinder efficient organization of different types of data within a task environment on-screen. Time can be a useful contextual cue to improve organization, as users often work with data most related to a task within a small window of time and have a sufficient memory of the relative history of their own interactions with data[1,2] Users face an ever-increasing amount of data available to them in forms such as work-related documents, personal and business email and IM, photos, music, web pages and other media and communications. Data across different types and from various applications are highly related to each other within the user's view of his or her tasks, yet it remains frustrating to manage disparate data through traditional hierarchical storage and access interfaces[3]. In this paper we suggest solutions to the problem of increasing complexity in relating and retrieving user data by adding time-based summary tools to the current desktop. The Desktop History Tool presents richer views on user activity, inactivity, frequently used or related documents, and the important items users care about.

Our goal in providing these summaries was to spur memories of related data and aid search and retrieval of important items. To accomplish this, we created a Desktop History Tool (or DHT), a prototype component on the desktop which maintains a temporal record of a user's interaction history and provides an interface to summarize and search that history for the relevant user data to a task. The primary goal of the DHT was to allow users to quickly recreate their previous task environments from the desktop, by augmenting the desktop metaphor rather than replacing it. We will discuss the functionality of the DHT, its design

elements, describe a user study we conducted to evaluate effectiveness, and discuss the results.

## Background

Users often express difficulty in quickly locating and retrieving data from disparate locations and sources to be able to use them together in a task at hand. We define these multi-document tasks as user *contexts*, or the complete set of data that might be helpful for a user to complete their task. For example, in the authoring of a research paper, the context spans email between authors, previous papers for formatting and reference, web pages on the relevant topic, on-line citations and related work, conference submission requirements, spreadsheets of collected data, a chart or graphing program and resultant image files, a slide-show editor for oral presentations, etc. The lifetime of this task is many weeks, and many of these elements will appear in windows on-screen, perhaps simultaneously, or perhaps within the same hour or day. The difficulties of restoring a desired subset of these elements vary by type: data that is not explicitly a file on local storage requires even more effort to retrieve through a specific network application interface such as an email application or a browser. Currently, users are called upon to explicitly organize their data to make later retrieval easier, by saving in categorized subfolders, customizing their bookmarks or inbox with named folders, etc. Many shortcuts have been devised such as 'Most-Recently-Used' lists of documents, either per application or across them, to alleviate some of the burden of navigation. Nevertheless, as users manage the increasing complexity of the set of items that comprise a project, some data that is peripheral to the task (and may need infrequent retrieval) is not well organized [4]. When such data is desired later it may have to be painstakingly restored.

## History as Context Restoration Tool

We propose that a user's own interaction history provides an invaluable self-organizing structure that future desktop user interfaces must recognize and incorporate. Such history already exists within some individual applications such as web browsers, an email folder of received mail viewed chronologically, or domain-specific activity logs such as the Outlook Journal, but a unified desktop history should prove far more valuable to users.

### Scenario for Context Restoration

A user found and then cut and pasted directions from several web sites for hotel accommodations into a wedding reminder email to family and friends. She later wishes to offer more choices from the initial search that she had omitted from the first email. We suggest that the clearest means to retrieve this web data in context starts with the fact that she knows that her web search occurred close to the time that the email was composed. Thus a helpful tool should allow browsing for and recall of both the email and the web data temporally, as temporal events are often suggested by users as locations around which to organize information, even when tools do not usually support this feature[5]. A prospective interface of this kind also allows for more awareness of peripheral data that on its own might not be explicitly recognized or valued. To complete the task currently, she may have to recreate the search that produced those extra pages, or go to the web history with fewer contextual cues. If she can rely on web history going back sufficiently, she must still be able to recognize the domains of the desired results before viewing the content, and discover them in a relative history consisting solely of other web sites. Even then she may need to scan specific urls within domains or spend time to partially re-navigate to them from the top level. And she still needs to retrieve the original email from the email inbox to access the right set of addresses for the update, before the context is sufficiently restored to complete the task.

### Related Work

The insufficiency of a desktop metaphor for data management is a well discussed and widely adopted view. And accordingly, time-based computing is an actively investigated solution space. There are many prototypes and research into time-based visualizations of computer usage for information retrieval. Some focus on the use of the world wide web but draw conclusions that are appropriate to all data types. Exemplars of this approach include the work of Wolber and Witchel et. al. called *WebTop*[6] which incorporates explicit user bookmarks, page history from usage, and the implicit content references within viewed pages to create a model for navigation in context. *InterLocus*, by Hayashi, Nomura, et. al. is a novel system for data management that implicitly places documents into a web-like storage and retrieval framework and hyperlinks documents into *temporally threaded workspaces*[7]. This work also applies to document state (undo, versioning, and multi-user

collaboration) and performs extensive interaction history collection and visualization. *InterLocus* depends on a desktop organized entirely via network access to the web that remains to be sufficiently well established to account for all useful data. The collection of user interactions with their data for future use is well discussed in related work by Bannon et. al.[8], Bauer[9], Curlander et. al.[10], Edwards et.al.[11], Derthick et. al.[12], and others. There are outstanding examples of interaction history visualization work that span all data types, including Hill and Hollan's *Vita Service*[13], and *LifeStreams* by Fertig, Freeman et. al.[14] In the *LifeStreams* architecture, a visualization of a users' stream becomes the central top-level organizing principal for their machine use, it essentially replaces the desktop, rather than augmenting it. This makes evaluation difficult across a diverse range of users and data. Another system for workspace visualizations through time is the *TimeScope* system by J. Rekimoto [15]. This work presents a flow of icons on a simulated desktop fluctuating as it is browsed through time, as a possibly more effective means of organizing and retrieving user data. While Rekimoto proposes future application interfaces would support navigation of this time-based space to implicitly restore the data that was being shown at any time, for the most part the user must still perform graphical maintenance on the desktop, adding and removing files, bookmarks etc. to group data, rather than visualizing the history of user interactions with data implicitly by use.

Overall, the requirements of novel infrastructure and the complexity of presenting time-based visualizations (from one to four dimensions) on the screen to users seem to be large factors inhibiting deployment and evaluation of time-based prototypes. Rather than seek means to reinvent a users' current UI, we have adopted a strategy of augmentation.

## DESKTOP HISTORY

We designed the DHT to assist users in finding the data that is important to them more efficiently and effectively than they are able to with current means. By using a relatively simple time-based summary of all of a user's activities—the tasks they complete and the temporal relationships between items (e.g., documents, URLs) associated with them—we hoped to capitalize on the rich contextual information that is coded by people as they work, information that is often lost in traditional storage methods.

We developed the DHT prototype for Windows XP, based on a traditional calendar and a journal output window that interoperated with a layer of instrumentation software responsible for collecting and storing the user's interaction history. The DHT calendar provided a quick overview of the interaction history, while a detailed activity summary in a lightweight journal format augmented the traditional desktop by allowing users to browse and launch any documents, web pages, and communications from a single entry

point organized by time. The DHT is shown below in a screen capture, Figure 1.

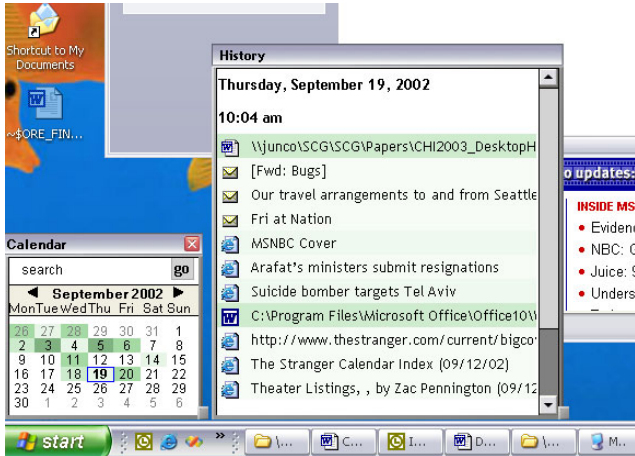


Figure 1. The Desktop History Tool (Version 1) on Windows XP

### Data Collection

In order to make the history of all interactions across any application available at one source, a data collection process was devised to collect a running summary of users' explicit interactions with any window. The data collection could be started or stopped at any time by the user according to 'stop' and 'start' shortcuts in the start menu. Collection included:

- Path and binary name of the executing program
- Time window was opened or became the input focus
- Title text of the window
- 'Handle' (System ID), Location/Size of the window
- Count of keystrokes and mouse movements within the window (though not the content of the input or output of any window)
- Path of documents created or accessed by window
- Time the window was closed or lost the input focus

Multiple sessions applied to each window through its lifetime. For example, if a user opened two windows in succession and then switched back and forth between them, a total of four recorded 'sessions' would be produced, switching back and forth again would produce 2 more, etc. Minimizing or maximizing a Window was treated by collection as a loss or gain of focus and thus completed or started a new session, respectively. Each window session was stored in the local database accessed by the DHT UI.

### UI Design: Visualization

DHT Design consisted of three principal elements: a *Calendar Window* for quick navigation and activity overview; a *History Window* for display of and navigation to specific user data, and a *Search Bar* to find, retrieve and summarize term-specific data.

### Calendar Window

A month-at-a-time calendar view was chosen as the top-level navigation for the DHT as it seemed the most intuitive way to display activity history, and selection of dates in a calendar grid is already well understood by most computer users. Search features were emphasized by placing a search text entry box at the top of the Calendar Window (or CW) with a 'go' button to execute search on any terms typed in. The CW was moveable and sizable but by default appeared in a small configuration to be less obtrusive and allow unfettered use of the desktop without taking focus away from daily tasks.

The CW displayed an overview of the user's total activity on the computer within a given month across individual days, showing also the current date and a selected day. It computed total activity by adding up the number of window sessions in each day that contained keystroke input or significant mouse movements and representing that total as a normalized gradient color on the day, as shown in figure 2. This colored gradient allowed users to quickly visualize and compare their relative activity between very active days (saturated color) and inactive ones (no color) in the smallest amount of space possible. Gradation also provided cues as to what overall usage activity matched up with any visible day and date to enhance recall and aid navigation[16]. A user could click a day in the calendar to select it and view detailed results for the day in the accompanying History Window.



Figure 2: Color gradient denoting relative user activity levels.

### History Window

The History Window (HW) displayed detailed summaries of user interactions on data for a given day, or on data that matched a given search term. In both versions produced, the results appeared as a scrollable list of data links in a straightforward journal format. Version 2 is shown below in Figure 3. In Version 2, actual meetings and calendar appointments from Microsoft Outlook were integrated within the HW to provide further context. This enabled users to look *forward* in the HW from the present day as well as back. The HW was laid out in the style of a journal, with the selected date at the top, and listed data below that by 1 hour increments. By browsing the HW a user could see what was worked on during their day as recorded in the collected sessions. Each entry of data in the HW was surrounded by a color gradient rectangle indicating the relative amount of keystroke and mouse activity by the user regarding that piece of data, within the hour increment. Thus, data worked on for more than an hour appeared in two consecutive increments. Any period of inactivity by the user of more than 5 minutes was suggested by a line of white space

between data listings. The last recorded activity in any hour was listed as a heading completing that entry for that hour: As in Figure 3, if a user worked on a spreadsheet until 12:17, and then left for lunch until past 1:00, 12:17 would cap off that hour's activity.

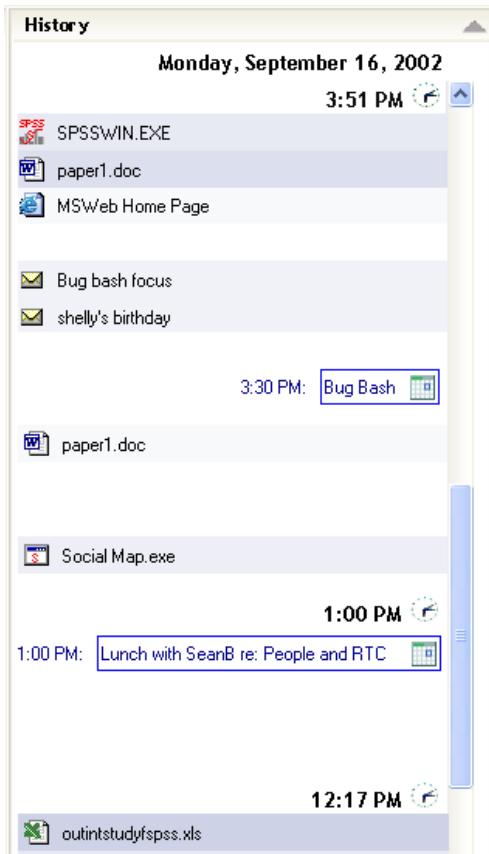


Figure 3: DHT History Window (V2) with meeting integration

Users could select any piece of data within the HW and double-click on it to launch and work with the data in the relevant application. Drag and Drop was also supported to allow items to be copied to storage, inserted in emails or sent to other applications.

#### Search Bar

The Search bar allowed partial or whole keyword search on all textual data recorded in data collection. Searching on specific terms would change the CW gradient color to orange to reflect gradation of activity results on window sessions that matched the search term. Executing a search would also color the 'go' button in the same color if good search results were available. The HW would initially summarize search results across all relevant days in the visible month (see Figure 4). Subsequently, clicking on a day in the VW in which good search results were indicated would narrow the summary of found data in the HW to only the selected day. Gradients on found data were drawn in a manner similar to the regular HW gradients--a total count of keystrokes and mouse movements for the data within an hour increment. Clicking on a day with no search results or clearing the search text box would reset the

search and remove the orange color replacing it with normal activity gradients again. In the HW, results produced by search were displayed and functioned in a fashion identical to standard HW data views for launching, drag and drop, etc.

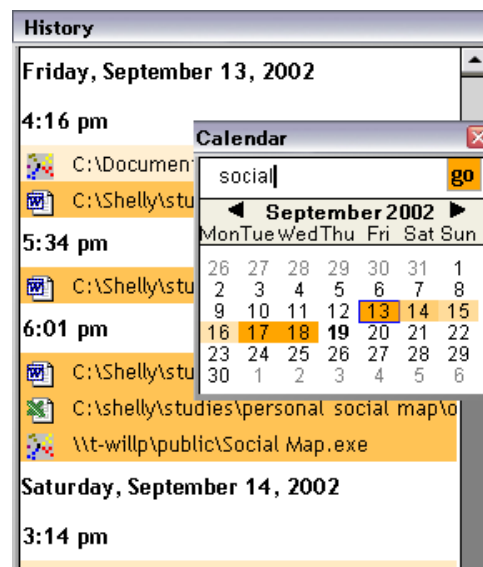


Figure 4. Search results in Calendar and History. Days in the Calendar with darker color gradient suggest higher frequency of items that contain the term 'social' with which users have been active. Data in History Window (V1) are colored by user activity on items which match the term.

#### USER STUDY

In an experimental study, participants first installed the data collector (described earlier) that gathered information about the tasks they performed and the data associated with those tasks. After 2 weeks of data collection, we presented users with the Desktop History Tool UI (V1) and asked them to explore its representation of their own interaction history. Participants then performed a series of tasks with their information that included recalling the important items they had worked on in the past week and recreating their computer workspace as it appeared in the morning of the day before. We also gathered information about the usability of the application (e.g., the search functionality) and information about how to increase its utility for users to guide future directions and features.

#### RESEARCH METHOD

##### Participants

Participants in this study were 13 employees of a technology company (9 females and 4 males) who volunteered to participate in return for a free cafeteria lunch. Participants were recruited to represent a variety of backgrounds and job experiences. Participants had been at the company an average of 4.1 years ( $SD = 4.7$ ). Participant users were required to install the small data collection component (described earlier) at least two weeks prior to the study date.

## **Procedure**

Each study session was conducted in the user's office and lasted approximately 1.5 hours. The experimenter observed the users as they completed the tasks, and all sessions were audio-taped for later analysis. Participants first completed several background questions such as job titles and the number of years they had worked at the company. They were also asked to estimate the percentage of work time they typically spend organizing the important data on their computers and give a subjective report of the amount of frustration they experienced at this task. After this they were asked to complete a series of tasks designed to evaluate the DHT interface and compare the DHT to users' own methods for finding important documents and data organization. At the conclusion of the session, users were debriefed about the goals of the study.

## **Tasks**

### *Listing Important Items (part 1)*

To test the efficacy of the DHT, we designed a series of tasks for users to complete using their normal computing methods or through the DHT. Users were first given 5 minutes to list the data they had worked on in the past week that they considered to be most significant to their work using any method they wished (A list item was defined as any document, email, web page, spreadsheet, graphic, picture or other data considered useful to their work). Users then ranked the top three items in terms of their importance, and gave reasons they listed each item (e.g., used frequently, shared with others, used for long periods of time). They were then asked to list related items to each of the important items they listed earlier; a related item was any data that users felt was in some way associated with an important item. If they could not think of any related data, they were told to leave this space blank.

### *Exploring the Desktop History Tool*

To give users some experience using the DHT, the first task simply asked them to explore the application in any way they desired for five minutes, and to 'think aloud' as they did so. They were then asked to explain in their own words what the application does. At this point users were also asked to rate the extent to which seeing the items presented in the DHT caused them to recall events or activities that they had previously forgotten.

### *Recreating a Workspace*

Next, users were asked to recreate their workspace as it appeared in the morning on the day before the experimental session. Specifically, they were asked to list all of the significant items they worked on in the first hour of this previous workday. Participants were first given 3 minutes to complete this task from memory, and they were then given an additional 3 minutes to redo this task using the DHT. While recreating a task context that existed only the morning before is a relatively straightforward task, we felt that the DHT would also serve as a useful memory aid in this respect, as well.

### *Listing Important Items (part 2)*

After completion of the previous tasks, participants were given an additional 5 minutes to redo the first part of this task again (listing important items), this time using the DHT. They then compared the two lists side by side and noted any differences between them.

We did not counterbalance the order of the methods participants used because we wanted to evaluate the DHT's performance specifically to augment the methods people typically use to locate important items.

### *Estimating Activity Level*

We designed the DHT to give users information about their workflow and how much they worked on different data in a day. In this task, users were asked to think back over the past work week to assess which day they were most active on their computers and which day they were least active, using any method they chose, and then the DHT. We were also interested in how users might make additional inferences about their own activities from history, so they were also asked to describe how DHT was used when obtaining this estimate.

### *Searching with the Desktop History Tool*

As described earlier, the DHT contained a simple partial keyword search feature that allowed users to find items in the collected text of the DHT session database. Results were displayed in the History Window, and the days on which search results were found turned an orange color. In the next task, users were asked to search for the last item on the important items list they had developed in the earlier task by using any term with which to identify it. Users were asked what information they thought would be searched through this mechanism. After the search was complete, users were asked whether the results matched what they expected from search.

### *Additional Qualitative Information*

After they had concluded the set of tasks, users answered a series of questions designed to gather additional information about the utility and usability of the DHT. Because gathering individual information about users' work habits may lead to issues about personal privacy, we also asked participants for information about how they felt about this issue.

## **Results and Discussion**

Due to scheduling and usage issues, not all users completed all tasks, and the degrees of freedom in the statistical data presented below reflect differences in the resulting sample sizes. Where appropriate, qualitative information aids in understanding the quantitative findings are also presented.

### *Organization*

We first asked participants to provide a self-assessment of how organized they considered themselves to be ("On the whole, when it comes to arranging items on your computer, how organized are you?") on a scale from 1 (*not at all*) to 7 (*very*). The average rating was 4.7 ( $SD = 1.5$ ), suggesting

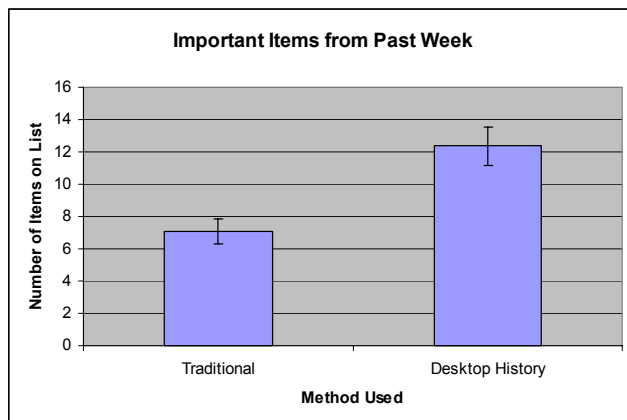
that the participants in this study considered themselves to be fairly organized overall. We also asked users to gauge how much work time they spend organizing items on their computers (1 = none; 7 = a great deal), and provide an estimate of the percentage of their work time they spend organizing their computer-based information. Users reported that they spend a considerable amount of time organizing, both in the scaled item ( $M = 3.1$ ,  $SD = 1.3$ ), and in the percentage estimate ( $M = 11.1\%$ ,  $SD = 5.6$ ). These results suggest that while the users in this study considered themselves to be organized, they do spend considerable time organizing items so they can find them again, so it appears that designing the DHT to augment current organization methods can show meaningful results.

### Remembering Important Items

One of our main hypotheses was that the DHT would help users find more of their important items than previous methods. Participants were asked to use their computers to find and list the items they had worked on in the five working days immediately prior to the study session day, and then to redo the same task using the DHT application.

Participants' normal methods for finding important items included manually browsing through folders in the file system and typing searches into the Windows search control panel. They also scanned their email inboxes and folders, their appointment calendars, and various homemade to-do lists. Many first attempted to list the overall tasks they had been working on to help them remember items that were associated with the tasks.

Users recalled and listed on average 7.1 ( $SD = 2.8$ ) important items using their traditional methods, while they recalled 12.4 ( $SD = 4.3$ ) items using the DHT. The increase in the number of items reported using the DHT was significant,  $t(12) = 8.24$ ,  $p < .001$ .



**Figure 5.** Number of items from past week listed as important using traditional methods or the Desktop History Tool. Note: error bars represent standard errors.

During the portion of the task where users explored the DHT, it was common to hear them comment that seeing the listed data prompted additional memories (e.g., “Ah, I totally forgot I did that yesterday!”). To gather a quantitative measure of the extent to which occurred, participants were

asked to rate the extent to which looking at items in the Desktop History Tool caused them to remember incidents or events they had previously forgotten (1 = not at all; 7 = a great deal). Users reported that seeing the list of temporally associated items did indeed prompt additional memories to a moderate extent ( $M = 4.1$ ,  $SD = 2.1$ ).

The desktop history's effect on improving users' memories was also evident in the comments they made during the task and in observations during the study sessions. The effects were often obvious:

*“This actually reminds me of something I have to do—write a note to somebody!”*

As we had hoped, users also saw the value of context as a memory cue. Seeing prominent items led users to recall additional, related items:

*“This is great to find things because it gives you a better context.”*

*“Seeing these things brings to mind a lot of stuff that is related to my main things.”*

One of the primary proposals of this research is that time is an important contextual and organizing variable. Observations of how users sought out information about their recent interactions with their important items shed additional light on this position. The methods participants used to locate the important items they had worked on recently almost always included seeking them in the time-based repositories they knew already existed. For example, when they searched for documents, users frequently went to most-recently-used documents list. When they knew they had visited important web sites, they looked at the address bar in Internet Explorer that shows in a drop-down list the most recently visited sites. The fact that people already use the limited time-based information that is currently available to them provides strong support for this principle. A more central and unified repository such as the DHT should add value.

### Recreating a Workspace

We asked participants to recreate from memory their computer workspaces as they appeared on the morning before the study session, and they then repeated this task using the DHT. Once again, users listed more items as appearing in their workspaces when they used the DHT ( $M = 4.4$ ,  $SD = 2.1$ ), compared to when they attempted the same task from memory ( $M = 8.4$ ,  $SD = 2.8$ ),  $t(11) = 7.27$ ,  $p < .001$ .

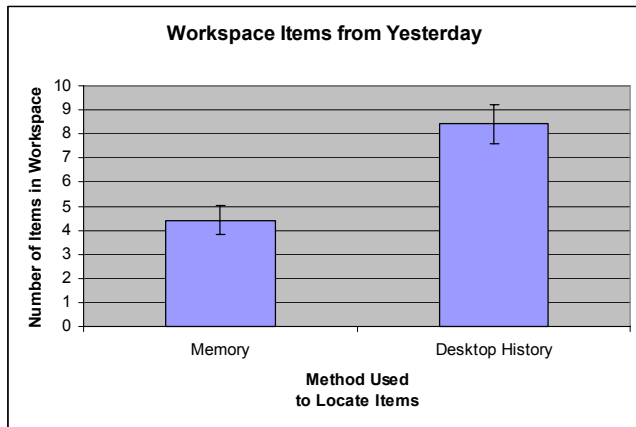
As they completed the tasks, users often voiced surprise at how difficult it was for them to recall information they had worked with only the day before. One user was unable to recall a single item that appeared in his workspace the previous morning. Overall, users tended to list very general items from memory:

*“I know I was working on email and some other documents, but I can't remember what they were specifically.”*

Users also had a strong tendency to list items that they knew generally occur on their prototypical workspaces, rather than listing specific items:

*“I know I usually do this in the morning, so I’m just going to put it down.”*

Participants found an obvious benefit when locating both important and peripheral data in the DHT.



**Figure 6.** Number of items recalled as appearing in the user’s workspace from memory or using the Desktop History Tool. Note: error bars represent standard errors.

#### Estimating Activity Level

The desktop history gives users information about their activity levels in two different ways. In the CW activity levels are suggested by the color gradients on days, and by hour in the HW. Assessing activity levels on data was also aided by the layout of the HW itself and the number of items listed. In the task that asked users to list their most and least active days, they often browsed the HW for the estimate, rather than the color saturation of the days in the CW. They typically selected each day throughout the week, browsed the HW and generally assessed their activity level by the number of items present or the size of the scroll bar. Some users mentioned that this visual comparison method might be somewhat inaccurate (even as they were using it) because each item in the HW was the same vertical height, and some data such as URLs were ones they only viewed only once or briefly in an hour.

#### Searching

Before users completed the search task, they were asked to describe what they thought the search set would be. All 13 users correctly assumed that the search mechanism would return results based on a words in the title of window as it appeared in the HW. When they were asked to search for an item on the ‘important items’ list they had written earlier, all 13 users were able to find the item using a key word search without prompting.

#### Effectiveness of Gradients

We were interested in whether users would spontaneously infer activity levels based on the color gradients in the CW and HW. While participants immediately knew that seeing color on a particular day meant they had been active on that

day, they did not at first infer that darker gradations necessarily meant more activity. The most likely reason for this appears to be that users were very active on nearly every day, making differences in the gradients less apparent. When the meaning of different color shades was explained to users, they immediately grasped its significance.

The meaning of the gradients in search was occasionally misunderstood by users. While all knew that the orange color on the calendar meant that data matching the search term had been found on the days colored orange, most (9/13) assumed that increased saturation of the orange color on some days meant that *more* data containing the search term had been located on those days, whereas gradients were designed to indicate the overall total of key and mouse activity on the found data, not the amount of results. Future work on the DHT will address how color saturation based on activity is calculated and understood.

#### Additional Qualitative Results

##### Overview or summary of personal data

One of the primary functions of the DHT is to give users a detailed history of their interactions with the data they care about. Observations of participants strongly suggested that they would find a great deal of value in more such high level or conceptual views of their information. Users also expressed a desire to select multiple days or non-standard time intervals in the HW in a manner similar to the multi-day summary available in search.

##### Information about inactivity

The version 1 HW that was shown to users in the study was solely information about the items they had interacted with on the computer. The periods of time when users were not at their computer showed up as white space in the HW. Users expressed a desire to have more information about periods of seeming *inactivity*. A commonly mentioned example of such information was electronically scheduled meetings. The mere presence of a meeting item in the desktop history may often be enough to evoke further memories about surrounding information. Overall, users thought that information about meetings was extremely valuable, and they felt it was natural to add meetings automatically in a calendar-type design. This information was added subsequently in DHT version 2.

##### Privacy and Sharing

The system we developed gathers information about a user’s computer-based activities, but the information is at present local and under the users ultimate control. Placing this data beyond their physical control raises issues about its privacy[13]. We also wanted to explore what users thought and felt about the privacy of their information. Overall, the reaction was somewhat positive. While some users expressed concerns about their information being gathered by ‘the wrong people,’ most were not concerned. Most participants thought that if the value provided by the application was great enough for users, they would probably not be very concerned about privacy in a single-user

setting. Many users expressed a desire to share their own explicitly selected views of the HW with others (e.g., a document or a URL) as a sharing mechanism. Several suggested that information about other users' activities in shared folders should be included in their history. This feedback suggests that privacy is not an overriding concern for users in a work context, when they can control what is shared explicitly.

#### *Uses for the Desktop History Tool*

Users found value in the application, as suggested by the fact that several of them asked to keep it on their machines and wanted to know where they could obtain updated versions (such as V2 with meeting integration) as they became available. In addition, its value was expressed in comments throughout the study such as:

*"This is actually a pretty cool thing!"*

*"I would definitely use this."*

Users thought of many different activities for which they could use the DHT apart from organization. Some thought it would make a good tracking tool, to allow them to see what types of things they do each day, how much time they spend on particular tasks, and discover ways to better organize their work time. Others thought it would be ideal as a tool to auto-generate information for periodic status reports or to include its output in project overviews.

#### **SUMMARY AND CONCLUSIONS**

The Desktop History Tool was created to give people a simple time-based representation of their interactions with data on their computers, integrated into the environment they already use. We found that the DHT successfully augmented the methods people typically used to find their important data. The DHT also enhanced people's ability to more fully reestablish their recent context to complete longer term tasks, or review shorter ones. We found that people already use temporal associations as memory and navigational aids wherever they can to find important data (e.g., 'most recently used' lists, appointments, etc.).

Based on the results of this study, we have redesigned the DHT in significant ways. These include forward-looking history that integrates with current appointment/meeting technology, detailed data previews that pre-cache and visualize the data within the History Window itself (rather than after launching), and encapsulation of data of multiple individual emails into threads for better summarization. Future versions of the DHT will also take into account that people often work across multiple machines (e.g., a desktop and a laptop), so information from each machine must be synchronized and integrated. Further scenarios on multi-user collaboration via shared history and using history as a sharing mechanism are also being explored. Accessible and highly integrated time-based visualizations of user interaction history should provide a valuable means for users to restore richer contexts and complete tasks more efficiently.

#### **Acknowledgements**

Many thanks to Lili Cheng, Shelly Farnham, David Vronay, Cezary Marcjan, Andrzej Turski, and Ryszard Kott for their contributions to this project.

#### **REFERENCES**

1. Freidman, WJ. and Wilkins, AJ. Scale Effects in memory for the time of past events. *Memory and Cognition* 13: 168-175.
2. Freidman, WJ. Memory for the time of past events. *Psychological Bulletin* 113 (1993), 44-66.
3. Cook, T. It's Ten O'Clock: Do you know where your data are? *Technology Review* (January 1995), MIT Press
4. Malone, TW. How do People Organize Their Desks? Implications for the Design of Office Information Systems. *ACM Transactions on Office Information Systems* 1 (1983), 99-112.
5. Gwizdka, J.: Timely Reminders: A Case Study of Temporal Guidance in PIM and Email Tools Usage. *Proceedings of CHI* (2000) 163—164
6. Wolber, D., Witchel, J., Kepe, M. and Ranitovic, I. Exposing Document Context in the Personal Web, *Proceedings of the International Conference on Intelligent User Interfaces* (IUI '01), ACM Press (online).
7. Hayashi, K., Nomura, T., Hazama, T., Takeoka, M. And Hashimoto, S. Temporally-Threaded Workspace: A Model for Providing Activity-based Perspectives on Document Spaces. *Proceedings of HyperText '98*, ACM Press (online).
8. Bannon, L., Cypher, A., Greenspan, S. and Monty, ML. Evaluation and Analysis of Users' Activity Organization, *Proceedings of CHI '83*, ACM Press (online).
9. Bauer, D. Personal Information Geographies. *Doctoral Consortium at CHI '02*, ACM Press (online).
10. Curlander, D. and Feiner, S. A History-Based Macro by Example System. *Proceedings of UIST '92*, ACM Press, 99-106.
11. Edwards, WK., Igarashi, T., LaMarca, A. and Mynatt, E. A Temporal Model for Multi-Level Undo and Redo. *Proceedings of UIST, '02*, ACM Press (online).
12. Derthick, M., & Roth, S. F. Data Exploration across Temporal Contexts. *Proceedings of IUI* (2000), New Orleans, LA. 60-67.
13. Hill, W. and Hollan, J. History-Enriched Digital Objects. *Conference on Computers, Freedom, and Privacy (CFP)* (1993),
14. Fertig, S., Freeman, E. and Gelertner, D. Lifestreams: An Alternative to the Desktop Metaphor, *Proceedings of CHI '96*, ACM Press (online).
15. Rekimoto, J. Time-machine Computing: A Time-centric Approach for the Information Environment. In *Proceedings of UIST* (1999), ACM Press, 45-54.
16. Gibbons, J.A. and Thompson, C.P. Using a Calendar in Event Dating. *Applied Cognitive Psychology* 15 (2001), 33-44