Data Mountain:
Using Spatial Memory for Document Management

George Robertson, Mary Czerwinski, Kevin Larson,
Daniel C. Robbins, David Thiel, and Maarten van Dantzich
Microsoft Research
One Microsoft Way
Redmond, WA 98052,USA
Tel: 1-425-703-1527
E-mail: ggr@microsoft.com

ABSTRACT

Effective management of documents on computers has been a central user interface problem for many years. One common approach involves using 2D spatial layouts of icons representing the documents, particularly for information workspace tasks. This approach takes advantage of human 2D spatial cognition. More recently, several 3D spatial layouts have engaged 3D spatial cognition capabilities. Some have attempted to use spatial memory in 3D virtual environments. However, there has been no proof to date that spatial memory works the same way in 3D virtual environments as it does in the real world. We describe a new technique for document management called the *Data Mountain*, which allows users to place documents at arbitrary positions on an inclined plane in a 3D desktop virtual environment using a simple 2D interaction technique. We discuss how the design evolved in response to user feedback. We also describe a user study that shows that the Data Mountain does take advantage of spatial memory. Our study shows that the Data Mountain has statistically reliable advantages over the Microsoft Internet Explorer Favorites mechanism for managing documents of interest in an information workspace.

KEYWORDS

3D user interfaces, desktop VR, information visualization, spatial cognition, spatial memory, document management

INTRODUCTION

Managing documents effectively on computers has been a key user interface design problem for the last thirty years. The issue has become more critical as users venture onto the World-Wide Web, because the number of easily accessible documents has increased dramatically. Graphics technology, processor speed, and primary memory capacity advances have made it possible to build systems that help with this document management problem.

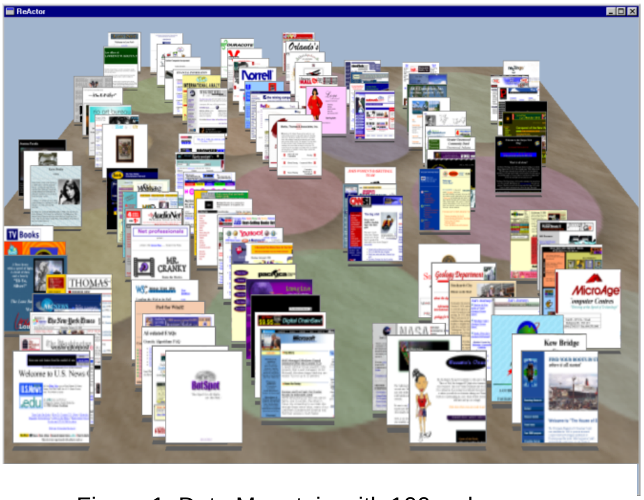

Figure 1: Data Mountain with 100 web pages.

The Data Mountain (Figure 1) is a novel user interface for document management designed specifically to take advantage of human spatial memory (i.e., the ability to remember where you put something). In our current prototype, the user freely arranges document thumbnails on an inclined plane textured with passive landmarks. We use 3D visual and audio cues to enhance the similarity to real-world object arrangement, yet use simple 2D interaction techniques and common pointing devices (like the mouse) for all interactions. The system is designed with a fixed viewpoint, so users need not navigate around the space. Users can identify and distinguish documents both through their thumbnail representation and through pop-up titles.

In this paper we describe the document management task and discuss existing graphical solutions. We then discuss related work from the field of spatial cognition, and issues of navigation and document management in the specific context of the World-Wide Web. Next, we describe the Data Mountain in detail, and report on a user study that compared it to the Microsoft Internet Explorer (IE4) Favorites mechanism. The paper concludes with a discussion of study findings and planned future work.

UIST '98. San Francisco, CA

Document Management

Document management tasks occur in a variety of contexts, and over a wide range of sizes of information stores and information structures. For example, tasks include managing files in a file system, mail messages, and web pages. The basic information structures a user might encounter include unordered sets, ordered lists, hierarchies, and graphs. Often the documents may belong to more than one of these information structures. Document browsing, searching, overviews, histories, and information workspaces all employ such structures.

The concept of information workspaces, introduced by Card, Robertson, and Mackinlay in 1991 [4], refers to the environment in which documents of interest can be compared, manipulated, and stored for future use. Iconic desktops, web browser Favorites or Bookmarks, and the Web Forager [5] are examples of information workspaces.

RELATED WORK

Document Management Systems

Early graphical methods for document management included list views, expandable lists for viewing hierarchies, and iconic 2D spatial layouts. The Apple Macintosh (circa 1984) included list views and a spatial layout (icon view). The spatial layout allowed the user to place icons in whatever grouping the user desired. Apple later added expandable lists for hierarchies, and piles [18]. Piles enrich the spatial layout by allowing the user to group related documents and take less screen space for the group. Yet another form of 2D spatial layout is the Treemap [14], which is a space-filling layout that is generated automatically, used primarily for overviews of document collections and their meta-data.

SemNet [10] was an early 3D spatial layout of documents. It tackled the difficult problem of visualizing networks. The result was difficult to understand because of the complexity of the information and layout. The Information Visualizer project [4][22] at Xerox PARC introduced a broad set of 3D visualization and interaction techniques for understanding information. In 1994, Maya Design Group introduced Workscape [3] as the first example of a 3D spatial layout of documents under the user's control. The Web Forager [5] built on the experience gained from the Information Visualizer project and introduced a 3D spatial layout for web pages and WebBooks [5].

The Visible Language Workshop of the MIT Media Lab did research in the design of dynamic virtual information spaces, combining typography and 3D graphics layout to present visually appealing interactive information landscapes [16][21].

Typical 2D desktops (Windows or Mac OS) use a spatial layout of icons and overlapping windows. As will be seen, the Data Mountain supports a larger number of objects, prevents overlap with a page avoidance algorithm, and uses thumbnail images instead of icons.

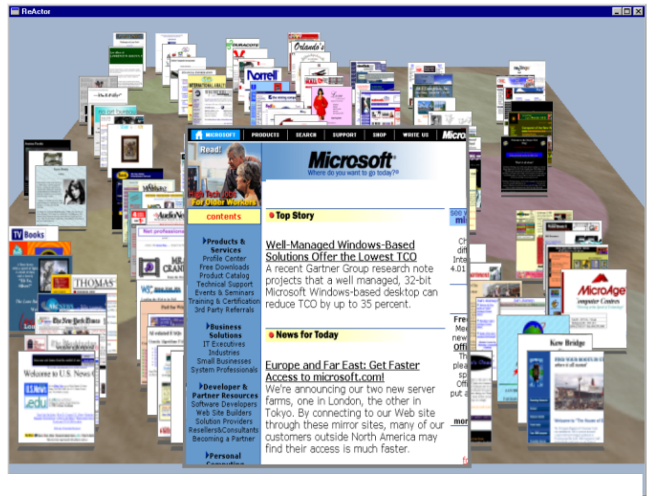

Figure 2: Selected page in preferred viewing position.

PadPrints [13] uses the Pad++ 2D zoomable user interface to implement a thumbnail image-based web history mechanism that is superior to text based history mechanisms. The authors raise a question about whether their technique is successful because of the thumbnail images or because of the zoomable interface. PadPrints uses an automatic layout for short term use; the Data Mountain uses a manual layout to exploit spatial memory for long term use.

Many of the techniques mentioned made use of spatial cognition, whether or not this was done intentionally. In particular, automatic spatial layouts of information leverage the user's ability to recognize and understand spatial relationships (both in 2D and 3D). The 3D interfaces make it possible to display more information without incurring additional cognitive load, because of pre-attentive processing of perspective views (i.e., smaller size indicates spatial relationships at a distance). The Maya Design Group Workscape system and the Web Forager intended to use spatial memory, by allowing the user to place documents as an aid to finding them again.

The Data Mountain is an advance over Workscape and Web Forager in several ways. First, the Data Mountain allows the user to place the document at an arbitrary location on its slope with a simpler interaction technique than the earlier systems (taking advantage of constrained motion along the inclined plane of the Mountain). Second, when a page is being moved, other pages are moved out of the way (active page avoidance), yet the user still sees visual cues indicating where every page will be when the movement is completed. Third, the Data Mountain exploits a variety of audio cues to augment the visual cues. Fourth, page titles are displayed whenever the mouse moves over a page. And fifth, in light of research in spatial cognition and wayfinding, visual neighborhood demarcation cues are provided to assist the user in arranging her personal space on the Data Mountain.

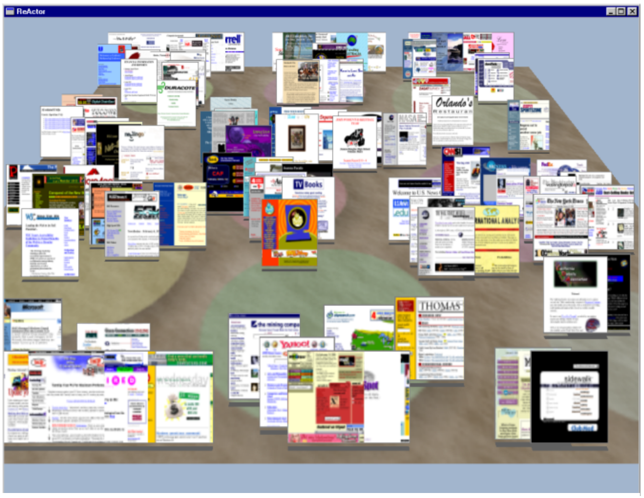

Figure 3: Second subject's layout of same 100 pages.

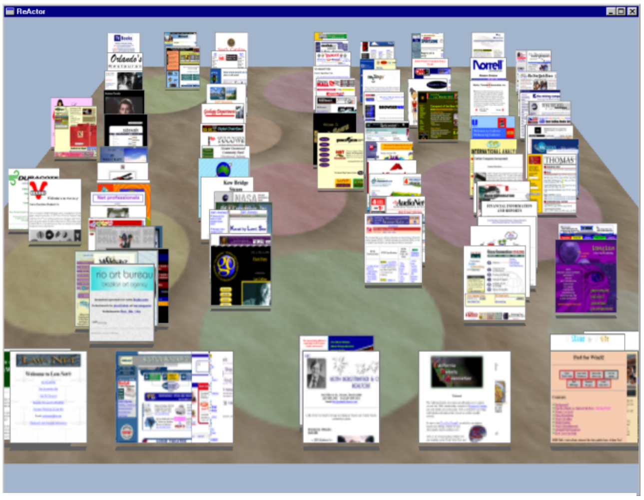

Figure 4: Third subject's layout of same 100 pages.

Spatial Cognition
There is a large body of literature on spatial cognition (see [23][7] for recent examples) and wayfinding [8][9][25], both for real and electronic worlds. Some of these studies have culminated in a set of guidelines for designers of virtual worlds. For instance, leveraging knowledge from the architectural domain [17][19], Darken and Silbert [9] have shown that adding real world landmarks, like borders, paths, boundaries and directional cues, can greatly benefit navigation performance in virtual reality. In their studies, Darken and Silbert have shown that stationary or predictably moving cues are optimal, and that multiple sensory modalities can be combined to assist searching through an electronic space (like 3D sound cues). They also have shown that if the space is not divided using a simple, organizing principle, users will impose their own, conceptual organization upon the space.

DATA MOUNTAIN

The Data Mountain is a 3D document management system. The current prototype is being used as an alternative to current web browser Favorites or Bookmark mechanisms, so we sometimes refer to pages as the objects that appear on a mountain. It should be understood that other forms of documents should work equally well.

When a page is first encountered, it appears in a preferred viewing position (see Figure 2), so that it is easily read. The user can place the page by dragging it with a traditional left-mouse-button drag technique. As the page is being dragged, other pages move out of the way so the page being moved is not occluded. After the page has been placed, it can be selected with a single click to bring it back to the preferred viewing position. Visual and audio cues, as well as the interaction techniques are more fully described below.

The Data Mountain is designed to work in a desktop 3D graphics virtual environment (also known as desktop VR), although it could certainly work in either Fishtank VR [27] or VR with head-tracked head-mounted displays. Examples described here are desktop VR examples.

Leveraging Natural Human Capabilities
The primary motivation for the design of the Data Mountain came from a desire to leverage natural human capabilities, particularly cognitive and perceptual skills. In particular, 3D perception is used to allow for the representation of a large number of web page thumbnails with minimal cognitive load. Our pre-attentive ability to recognize spatial relationships based on simple 3D depth cues (like perspective views and occlusion) makes it possible to place pages at a distance (thereby using less screen space) and understand their spatial relationships without thinking about it. We can leverage audio perception to reinforce what is happening in the visual channel. Both the visual and auditory perception can enable basic human pattern recognition capabilities. And finally, we hope to use spatial memory to make it easier to find documents in the information workspace.

In the real world, spatial memory often aids us in finding things. For example, when we place a piece of paper on a pile in our office, we are likely to remember approximately where that paper is for a long time. Our hope is that this ability works as well in a virtual space as it does in the physical world. This is not an obvious conclusion. However, if spatial memory is primarily an act of building a mental map of the space, then we should be able to do the same thing in a virtual environment, and take advantage of it.

Data Mountain Visual and Audio Design
The Data Mountain provides a continuous surface on which documents are dragged. The document being dragged remains visible so the user is always aware of the surrounding pages. This is in direct contrast to the way a Web Forager [5] user places documents in a discrete set of tiered locations using flicking gestures. We believe the

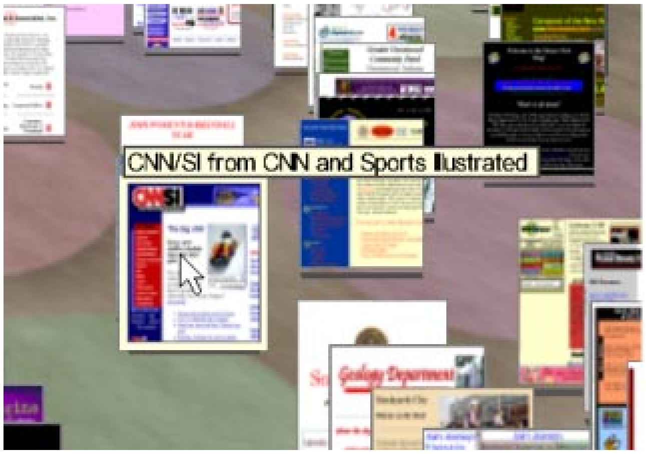

Figure 5: Title shown while hovering over page.

user's act of directly placing the page on the continuous surface of the Data Mountain aids spatial memory.

The Data Mountain prototype uses a planar surface (a plane tilted at 65 degrees), as shown in Figure 1. The landscape texture on the Data Mountain surface provides passive landmarks for the user meant as an aid for grouping objects into categories, but the landmarks have no explicit meaning. The user can place the web pages anywhere on the mountain. In practice, users create meaning by organizing the space. In our study, there were many ways to lay out the same set of pages (compare Figures 1, 3, and 4). There is no "right" layout; rather, the layout is very personal, has meaning for the individual who created it, and evolves over time under user control.

Note that the current prototype provides no mechanism for labeling groups of pages with category titles. While some users requested this feature, we found they built a very accurate mental map of their categories even without explicit labels. Some users employed particularly salient thumbnails as visual identifiers of their groups, keeping them in front of all other members of the group (thus creating their own landmarks).

There are a number of 3D depth cues designed to facilitate spatial cognition. The most obvious are the perspective view and occlusion, particularly when pages are being moved. The landmarks also offer an obvious cue, which may or may not be utilized during page placement as well as retrieval. Less obvious, but also quite important, are the shadows cast by the web pages.

Subtle but pervasive audio cues accompany all animations and user actions to reinforce the visual cues. The sound effects are highly dynamic. For example while moving a page the user hears a humming sound that changes pitch based on the speed of the page as it is dragged, as well as indicating spatial location by controlling volume, low pass filtering, panning, and reverb level. As the user moves a page, other pages move out of the way as needed, producing another distinctive sound.

Data Mountain Interaction Design

When the user clicks on a page stored on the Data Mountain, the page is moved forward to a preferred viewing position, as shown in Figure 2. The animation to bring the page forward lasts about one second [4], uses a slow-in/slow-out animation [6][11], and is accompanied by an audio cue. We use a higher resolution texture map for the page image in the preferred viewing position, ensuring that the page is quite readable.

When in the preferred viewing position, a click on the page will either select and follow a hyperlink, or put the page back on the Data Mountain in its last known location. This is also done with a one second, slow-in/slow-out animation accompanied by audio.

We provided a pop-up label similar to tool-tips to display page titles. Subjects tended to use their spatial memory to get to the neighborhood of the page, then riffle through the titles (like riffling through a pile of papers on your desk) to find the page. A standard tool-tip uses a hover time before the tip is displayed. We determined in a pilot study that the hover time was not effective since it precluded rapid inspection of multiple titles. Hence, the title appears as soon as the mouse moves over a page. An example is shown in Figure 5. In one group from our user study, the title was shown just above (but disconnected from) the page. Some subjects could not easily distinguish which thumbnail the title applied to, so our second Data Mountain design added an identically colored halo around the thumbnail, creating a visual link to the title.

A page can be moved at any time by dragging it with the mouse. Since the page is visible during the move, the user knows where the page will be when the drag is terminated. The movement is continuous and constrained to the surface of the Data Mountain. This results in one of the principal advantages of the Data Mountain; the user gets the advantages of a 3D environment (better use of space, spatial relationships perceived at low cognitive overhead, etc.), but interacts with it using a simple 2D interaction technique.

Page Avoidance Behavior

When moving a page, what is the right behavior for pages that are encountered (i.e., how are collisions handled)? We have tried three alternatives, each improvement driven by user comments. First, we did nothing. That is, the page in motion simply passed through the pages in the way. This approach suffers from several problems. It makes the system seem lifeless and makes the metaphor harder to understand. In addition, it is quite easy to put pages right on top of each other, making it difficult to find some pages.

Second, we tried a simulation of tall grass. Think about what happens when you walk past tall grass. If you walk slowly, the grass moves out of your way slowly then returns. If you walk fast, the grass seems to fly out of your way. We implemented a simple simulation in which previously placed pages behave like grass displaced with a page the user is dragging. Each page that is moved out of the way uses a ½ second animation to move aside, followed by a one second animation to move back. This feels very lively but suffers from two problems. It still does not solve the problem that two pages can end up in the same location: the user may drop a page close to another document's real location while that document is temporarily displaced, soon to return, causing one to occlude the other. The collision avoidance behavior should be designed to eliminate such surprises. Also, the movement is based on first encounter. In other words, if after triggering avoidance the dragged page slows down or hovers in-place, the return animation will still take place, causing the objects to intersect anyway. Essentially, the 'grassy technique' works well for continuous dragging but tends to be annoying when you slow down. Note that slowing down is an essential part of object placement: effectively, the grassy technique caused some users to create unnecessary occlusion since they would pick the just vacated spot of a displaced page for the new page. It could be modified to be dependent on your speed; but since it is basically an estimate, it will fail in some cases. The first group of Data Mountain users described in the experiment below used this page avoidance mechanism. These problems provided strong motivation for finding another method and running another set of subjects, as some users effectively lost many pages due to occlusion.

In our current implementation, we continually maintain a minimum distance between all pages, even while a page is being moved, and transitively propagate displacement to neighbors as necessary. This has the advantage that the user dragging the page continually sees what state will result when the drag is terminated (i.e., there is no animation settling time). Also, the pages never get fully obscured. In particular, you cannot move a page and leave two pages in the same location. On the other hand, displacements may propagate far afield when a cluster of closely packed pages is 'pushed' by a dragged page, resulting in more visual unrest than is really desirable. Still, the approach feels quite lively, and was used by the final group of subjects in our study. We feel this change contributed most to improved user performance in the second Data Mountain group.

Implementation

The Data Mountain prototype runs under Windows NT version 4 on PCs equipped with Intergraph Intense 3D Pro 1000 or Pro 2200 graphics accelerators. All application code was written in C++, utilizing our own libraries for animation and scene-graph management. These libraries in turn used the ReActor [2] infrastructure and OpenGL as the underlying graphics library.

The interactive sound for the Data Mountain was based on MISS (the Microsoft Interactive Sound Sequencer) which takes parametric sound events and sequences them using MIDI to communicate to the wavetable synthesizer on a Creative Labs AWE64 Gold card.

For prototyping and user study purposes, the web pages in the current implementation are not "live", i.e.; one cannot select and follow a hyperlink. Future versions of the Data Mountain will contain live web browsing capabilities.

The 100 pages used in the study below are screen snapshots of actual web pages in 24-bit color. We employ two bitmap sizes of each page for texture mapping: a small 64x64 pixel version (12KB each) for the thumbnails on the Data Mountain surface, and a 512x512 pixel version (768 KB each) for the close-up view. One hundred thumbnails plus a close-up together will fit in 2MB of texture cache. Our system implements text labels using texture-mapped fonts [12]. This is vastly preferable to vector fonts, and has the advantage that it enables display of legible text on surfaces that are not screen-aligned. Category labels placed on the Data Mountain surface are just one potential use of such perspective-distorted text.

If pop-up labels are naively attached to thumbnails, they will be subject to perspective projection, and thus be smaller for pages that are placed towards the back of the information workspace. We found it difficult to choose an absolute label size in model coordinates that produces appropriately scaled text for labels of both foreground and background documents. Instead, we implemented labels for thumbnails to be of identical size, independent of the document's distance from the viewer in virtual space. We do this by placing the title-tip a constant distance away from the eye-point, on the vector from the eye-point to the page whose title we are showing.

USER STUDY

Studies by Tauscher and Greenberg [24] and Abrams [1] were the earliest attempts to gather information about user behavior as they traverse the web over several months. According to Abrams [1], users develop their own personal web information spaces through the use of Favorites mechanisms in order to combat the problems of information overload, pollution, entropy, structure and lack of a global view of the web. Users do this by building a smaller, more valuable, organized and personal view of the web.

Usage tracking shows that hotlists, bookmarks and Favorites folders are the navigation tools most frequently utilized by users for locating information on the web [20].

Hence, web browser designers need to provide their users with mechanisms for creating personal web information spaces that can reliably and efficiently return the user to their favorite web sites. Implementing such mechanisms relaxes the cognitive and temporal demands of hypertext navigation [1]. Usability studies, as well as basic research, however, indicate that the current designs for navigating the web are still sub-optimal in supporting users' cognitive models of web spaces and the amount of information they need to repeatedly consume [1][24].

The Abrams study [1], in particular, pointed out how episodic memory [26], or memory for events, could be

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.

WHAT WILL YOU BUILD? | **sales@docketalarm.com** | **1-866-77-FASTCASE**

fastcase® Smarter legal research.