

The Perspective Wall: Detail and Context Smoothly Integrated

J. D. Mackinlay, G. G. Robertson and S. K. Card

UIR-R-1991-03

THE PERSPECTIVE WALL: DETAIL AND CONTEXT SMOOTHLY INTEGRATED

Jock D. Mackinlay, George G. Robertson, and Stuart K. Card

Xerox Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304
415-494-4335, mackinlay.parc@xerox.com

ABSTRACT

Tasks that involve large information spaces overwhelm workspaces that do not support efficient use of space and time. For example, case studies indicate that information often contains linear components, which can result in 2D layouts with wide, inefficient aspect ratios. This paper describes a technique called the *Perspective Wall* for visualizing linear information by smoothly integrating detailed and contextual views. It uses hardware support for 3D interactive animation to fold wide 2D layouts into intuitive 3D visualizations that have a center panel for detail and two perspective panels for context. The resulting visualization supports efficient use of space and time.

KEYWORDS: User-Interface Design Issues - *visual output strategies, interface metaphors, graphic presentations, screen layout*; Analysis Methods - *analysis of contents of particular domains*; Domain Specific Designs - *information retrieval*

INTRODUCTION

Tasks that involve large information spaces require workspaces based on explicit techniques for handling the volume of information. There exist two basic strategies for developing such techniques: a *space strategy* uses layout and graphic design to pack appropriate information in one view, while a *time strategy* uses view transitions to spread information over multiple views.

Although the space-time distinction is useful, large information spaces often require simultaneous use of both strategies, which can involve interactions among the techniques. For example, case studies indicate that tasks often involve *spanning* properties (such as time) that structure the information linearly. This linear structure results in 2D layouts with wide aspect ratios that are difficult to accommodate in a single view. Current computer screens are quite small compared to "real-world" workspaces, such as dining room tables[3]. Furthermore, all workspaces are limited by human size and perception in the amount of information they can make visible in a single view. Layouts with wide aspect ratios must be reduced in scale to fit in a single view, which causes very small details. When a time strategy is used to enlarge details by only viewing part of the layout, the view may omit contextual information that is needed to work on a task or navigate to another part of the information space. Reintroducing the necessary contextual information interacts with the space strategy by reducing the area available for viewing details.

Resnikoff observes that the human eye and other biological systems process the vast amounts of information available in the real world by smoothly integrating a focused view for the detail with a general view for the context[10]. In particular, the retina of the human eye is hierarchically decomposed into a foveal region that perceives details and a surrounding low resolution region for daylight color vision and nighttime monochrome vision. Furthermore, other biological information processing systems, including bat echo-location, have similar decompositions. He argues that this architecture is a result of general information processing principles such as selective omission of information and aggregation into more abstract forms.

This paper describes a technique called the *Perspective Wall* that integrates detailed and contextual views to support the visualization of linearly structured in-

formation spaces. Companion papers describe our related visualization for hierarchically structured information spaces[12] and our general approach for developing a system that supports general information visualization[1]. In this paper, we begin by showing that work practice often causes information to have a linear temporal structure. Such structures result in 2D layouts with wide aspect ratios that are not adequately handled by existing techniques. The Perspective Wall technique takes advantage of hardware support for 3D interactive animation to imitate the architecture of the eye system. It folds a 2D layout into a 3D wall that smoothly integrates a region for viewing details with perspective regions for viewing context. This intuitive distortion of the layout provides efficient space utilization and allows smooth transitions of views. Analysis of the Perspective Wall technique indicates a threefold improvement over simple 2D visualizations. A version of the Perspective Wall has been implemented and incorporated into the *Information Visualizer*, a 3D Rooms system for visualizing information[11]. The *Information Visualizer* also includes support for multiple asynchronous agents and smooth interactive animation[11], 3D navigation[7, 8] and object manipulation[8]. The Perspective Wall has been used to visualize various types of temporally structured information, including a file system and corporate memoranda.

LINEAR STRUCTURE IN DOCUMENT GROUPS

Work practice often causes information spaces to be linearly structured by some metric, for example: project records (chronology) or directories (alphabetical order). The location of an item along this metric can be an excellent retrieval cue, either directly or in relation to "landmark" items whose location may already be known. For example, an architect whose office files we were studying observed that the titles of documents in his office were not good retrieval cues, but that creation time was. We thus need to be able to display linear structures in computational visualizations.

To better understand this linear structuring, we examined the documents from three actual projects. Two of these cases, Case 3-1075 and Case 3-1078, were from the files of the architect above, who serves as an expert witness for litigation concerning building construction. The files are comprised of construction and litigation documents for construction projects. The third case, Case P70/PPM, consisted of documents produced by a National Research Council committee in a study of human performance models suitable for computer-aided engineering of cockpit design. They include administrative, informational, and draft documents generated in service of producing a committee report. This case was chosen as a contrast to the architectural cases.

<i>Forensic architecture (Cases 3-1075, 3-1078)</i>	
Number of Documents: 431	
TASK CYCLE DOCUMENTS	
1.	<i>Building cycle</i>
	Contractual agreements
	Building plans
	Analyses
	Financial invoices
2.	<i>Litigation cycle</i>
	Complaints
	Court procedures
	Remediation documents
COORDINATION DOCUMENTS	
	Information requesting and giving
	Direct coordination
	Conflict resolution
<i>Study committee (Case P70/PPM)</i>	
Number of Documents: 224	
TASK CYCLE DOCUMENTS	
	Information gathering
	Analysis
	Report drafts
COORDINATION DOCUMENTS	
	Information requesting and giving
	Direct coordination
	Conflict resolution

Figure 1: Case study document categories.

Although the set of documents in each of these cases was not large (on the order of a file drawer), they were extensive enough that the users expressed frustration at the effort required to manage them and find information in them for their work. As a way of analyzing the collection, the documents were grouped into categories (see Figure 1). The resulting categories suggest that the documents have a simple overall structure reflecting processes that naturally develop through time. In the forensic architecture cases, there are two main task cycles: (1) construction of the building and (2) litigation over its defects. In the committee work case, there is one main task cycle leading to the committee report.

Figure 2 contains two plots that show the number of documents (with identifiable creation dates) as a function of time for construction litigation and committee work. These plots reveal further time-based texture in the document collections. The upper plot shows the documents of both forensic architecture cases plotted together, one in light bars and the other in dark bars. The construction and litigation cycles are clearly evident. The upper plot also indicates that information can have a layered structure caused by simultaneous tasks. The two litigation cycles were completely overlapped, and the original build cycles were partially overlapped.




Figure 2: The upper plot shows the number of construction litigation documents by quarter. The dark bars are Case 3-1075 and the light bars are Case 3-1078. The lower plot shows the number of committee work documents by month for Case P70/PPM.

lapped. We will show that the Perspective Wall can be used to visualize this layering of information in addition to its linear structure. The lower plot clearly shows the pulsing effect of the study committee meetings. Each meeting generated a flurry of preparatory documents as participants were forced to get their work done before the meeting deadline, to arrange agendas, and to prepare the meeting minutes.

These case studies suggest that spanning properties, such as time, often reveal some of the structure in an underlying document collection. The visualization of the document collection in terms of these properties can reveal subtle textures reflecting the underlying processes that generated the documents. These textures could be useful to the participants attempting to retrieve the documents because they could use their contextual knowledge of the underlying work process to aid the search. Linear visualizations based on time or some other spanning property could also allow non-participants to rapidly understand the collection as a whole and to search it.

VISUALIZING LARGE INFORMATION SPACES

The principal obstacles to a visualization of linear information structures are (1) the large amount of information that must be displayed and (2) the difficulty of accommodating the extreme aspect ratio of the linear structure on the screen.

Window systems provide some support for processing large amounts of information. For example, windows can be used as a space technique to group related information and as a time technique for switching among views (for example, overlapped windows, scrolling, or icons). Window systems can also provide some contextual information through the proximity of windows on the display. However, virtual desktops quickly become cluttered and unusable during work with large amounts of information.

The Rooms window manager is a time technique that increases the effective working size of a window system by allowing users to switch among window "working sets" as they switch tasks[3]. Rooms also supports navigation with various techniques including an "Overview Room" that allows users to see and work with the entire workspace. However, Rooms does not address the problem of developing visualizations for tasks involving large information spaces (except by moving the clutter from other tasks to other rooms).

A common technique for integrating detail and context is to have two simultaneous views: an overview with a scale-reduced version of a workspace, and a detailed view into the workspace where work can be accomplished[3, 4, 6]. The overview typically contains an indication of the detailed view's location that can be manipulated for rapid movement through the

workspace. However, a uniform scale reduction of the workspace causes it to appear very small. Furthermore, important contextual information, such as the neighborhood of the viewing region, is just as small as unimportant details. Finally, if the display space for the overview is increased to make the workspace appear larger, the space for the working view becomes too small.

Rather than a uniform overview of a workspace, an effective strategy is to distort the view so that details and context are integrated. Furnas developed a general framework called Fisheye views for generating distorted views[5]. Fisheye views are generated by Degree of Interest functions that are thresholded to determine the contents of the display. However, thresholding causes the visualization to have gaps that might be confusing or difficult to repair. Furthermore, gaps can make it difficult to change the view. The desired destination might be in one of the gaps, or the transition from one view to another might be confusing as familiar parts of the visualization suddenly disappear into gaps.

Spence and Apperley developed an early system called the Bifocal Display that integrates detail and context through another distorted view[9]. This 2D design is a conceptual ancestor of the Perspective Wall system described in this paper. The Bifocal Display was designed for professional offices that contain information subdivided into a hierarchy of journals, volumes, issues and articles. Abstractly, the workspace consists of information items positioned in a horizontal strip. The display is a combination of a detailed view of the strip and two distorted views, where items on either side of the detailed view are distorted horizontally into narrow vertical strips. For example, the detailed view might contain a page from a journal and the distorted view might contain the years for various issues of the journal. Because Bifocal Displays are two dimensional, they do not integrate detail and context completely smoothly or intuitively. Two versions of an item are required, one for the detailed view and one for the distorted view. The relationship between these versions may not be obvious. As the focus moves, items suddenly expand or shrink, which may be confusing. Furthermore, the distorted view treats all contextual items identically, even those near the detailed view.

PERSPECTIVE WALL

The Perspective Wall technique uses recent advances in hardware support for interactive 3D animation to address the integration problems of the Bifocal Display. A physical metaphor of folding is used to distort an arbitrary 2D layout into a 3D visualization (the wall), while automatically retaining any 2D task-specific features. More importantly, no special large and small scale versions of items must be designed (as in the Bi-

focal Display). The wall has a panel in the center for viewing details and two perspective panels on either side for viewing context (see Mackinlay Plate 1). The perspective panels are also shaded to enhance the perception of 3D. This intuitive visualization provides efficient space utilization for 2D layouts with wide aspect ratios. In addition, the vertical dimension of the wall can be used to visualize layering in an information space. The Perspective Wall in Mackinlay Plate 1 holds cards that represent files in a computer system that are structured by modification date (horizontally) and file type (vertically). The perspective view has the further advantage that it makes the neighborhood of the detailed view larger than more distant parts of the contextual view.

Smooth transitions among views is accomplished by allowing the user to adjust the wall as if it were a sheet in a player piano moving selected notes to the center of view. Mackinlay Plate 2 shows the visualization after the user has selected a card in the left perspective view of Mackinlay Plate 1.

The space utilization of a Perspective Wall can be analyzed by comparing it with the corresponding unfolded wall, which can be placed in the field of view to simulate a 2D visualization of the 2D layout. Figure 3 shows a top view of two different placements of the flat wall: (1) the flat wall completely in the field of view with small details, and (2) the flat wall at the same distance as the Perspective Wall detail panel with much of the layout outside the field of view. The analysis uses the following parameters: the angle θ_f , which is the amount of folding of the perspective panels, the width w , which is the size of a perspective panel (assuming a unit wide center panel), and the angle θ_e , which is half of the visual angle of the field of view. Typical values for these parameters are: $\theta_f = 60^\circ$, $w = 5$, and $\theta_e = 17.5^\circ$.

When the flat wall is placed completely in the field of view, the details are much smaller than the details on the Perspective Wall. The relative sizes of the details can be determined by comparing their distances to the eye (see the top analysis in Figure 3). The distance from the eye to the distant flat wall is $(\frac{w+0.5}{\sin(\theta_e)}) \cos(\theta_e)$, which comes to 17.4 for the typical values. The distance to the center panel of the Perspective Wall is $(\frac{w \cos(\theta_f)+0.5}{\sin(\theta_e)}) \cos(\theta_e) - w \sin(\theta_f)$, which comes to 5.2. Therefore, for typical values of the parameters, the details on the center panel are at least three times larger than the details on a flat wall that fits in the field of view.

When the flat wall is placed at the same distance as the Perspective Wall center panel, much of this closer flat wall is outside the field of view (see the bottom analysis in Figure 3). The amount of the Perspective Wall




Figure 3: A top view of the Perspective Wall (bold line) and two flat walls (hollow lines) positioned in the field of view (dotted line). The "distant" flat wall is placed to the right to fit in the field of view and the "detailed" flat wall is placed at the same distance as the Perspective Wall detail panel. The distant flat wall is analyzed in the top half of the diagram and the detailed flat wall in the bottom half.

in the field of view is, of course, $2w + 1$, which is 10.0 for the typical values. The amount of the closer flat wall is $2(w \cos(\theta_f) - (\frac{w \sin(\theta_f)}{\cos(\theta_e)}) \sin(\theta_e)) + 1$, which is 3.3. Therefore, for typical values of the parameters, the Perspective Wall makes three times as much information visible as a flat wall that has details of the same size.

A major advantage of the Perspective Wall is that its intuitive 3D metaphor for distorting 2D layouts allows smooth transitions among views. When the user selects an item, the wall moves that item to the center panel with a smooth animation. This animation helps the user perceive object constancy, which shifts to the perceptual system work that would otherwise have been required of the cognitive system to reassimilate the view after it had changed[11]. Furthermore, the relationship between the items in the detail and context is obvious. Items even bend around the corner.

The Perspective Wall has the additional feature that it is easy to allow the user to adjust the ratio of detail and context. This is quite important when the detailed view contains a lot of information. The metaphor is to stretch the wall like a sheet of rubber as shown in Mackinlay Place 3, which corresponds to the wall in Mackinlay Plate 2.

IMPLEMENTATION

The Perspective Wall technique has been implemented as part of the Information Visualizer, a system that uses 3D visualizations and interactive animation to work with large information spaces, especially when the structure of the information can be exploited[1, 11, 12]. A 3D version of Rooms[3] is used to support task switching. An animation loop[11] is used to support view transitions.

The Perspective Wall works with any 2D layout that has been described as a list of 2D vectors and 2D positioned text. On each animation cycle, three passes are made over this list to generate the three panels of the Perspective Wall. The placement of the 2D layout on the panels is determined by a single parameter that specifies what part of the layout should be in the center of the detail panel. The wall scrolls when this parameter is set to a new value, for example, when a card is selected. A governor is used to monitor the animation rate and adjust this parameter to create smooth scrolling[11].

The Perspective Wall has been used to visualize various types of information. The Mackinlay Plates represent files in a file system that are classified by their modification date and file type. Vacations and other work patterns are clearly visible. The technique has also been used for corporate memoranda and reports, which also have a useful linear structure. The technique is particularly effective when combined with a retrieval technique that allows the user to select an item and find similar related items[2]. The Perspective Wall makes it easy to visualize the results of such retrievals because it shows all similar items simultaneously and in context.

CONCLUSION


Case studies indicate that information spaces often have spanning properties such as time that lead to 2D layouts with wide aspect ratios. The Perspective Wall technique uses hardware support for 3D interactive animation to turn these 2D layouts into 3D visualizations that imitate the human eye and smoothly integrate detail and context. The technique uses space efficiently and allows the view to be adjusted smoothly and intuitively. The 3D metaphor also allows the ratio of detail and context to be smoothly adjusted.

The Perspective Wall is part of a larger effort to develop techniques for accessing and managing large information spaces and thus to support large scale cognition[1]. Our approach is to exploit human experiences and capabilities by using emerging technologies for 3D visualization and interactive animation to develop highly interactive user interfaces. Our experience so far suggests that it is possible to create effective space and time techniques that improve management and access of such information spaces. Interactive animation, in particu-


lar, seems to be a powerful technique because it shifts cognitive load to the perceptual system and thus supports smooth transitions among views. We have been able to use these techniques to visualize the structures of information spaces, such as entire file systems, that have never been seen before. Our initial prototypes suggest that highly interactive user interfaces are likely to support large scale cognition and thus deserve further research activity.

References


- [1] Card, S. K., Robertson, G. G., & Mackinlay, J. D. The Information Visualizer: An Information Workspace. *Proceedings of CHI'91* (New Orleans, Louisiana).
- [2] Cutting, D. R., & Pedersen, J. O. Optimizations for dynamic inverted index maintenance. *Proceedings of SIGIR'90*, Brussels, Belgium, ACM Press, September, 1990.
- [3] Henderson, D. A. Jr., & Card, S. K. Rooms: The use of multiple virtual workspaces to reduce space contention in a window-based graphical user interface. *ACM Transactions on Graphics*, 5(3, July), 211-243, 1986.
- [4] Fairchild, K. M., Poltrock, S. E. & Furnas, G. W. SemNet: Three-dimensional graphic representation of large knowledge bases. *Cognitive Science and its Application for Human-Computer Interface* R. Guindon (Editor). Lawrence Erlbaum, New Jersey, 1988.
- [5] Furnas, G. W. Generalized fisheye views. *Proceedings of CHI'86* (Boston, Mass), 16-23.
- [6] Herot, C. F. Spatial management of data. *ACM Transactions on Database Systems* 5 (4, December), 493-514, 1980
- [7] Mackinlay, J. D., Card, S. K., & Robertson, G. G. Rapid controlled movement through a virtual 3d workspace. *SIGGRAPH '90 Conference Proceedings* (Dallas, Texas, August 1990). *Computer Graphics*, 24, 4 (August 1990), 171-176.
- [8] Mackinlay, J. D., Card, S. K., & Robertson, G. G. A semantic analysis of the design space of input devices. *Human-Computer Interaction*, 5, 2-3, 145-190, 1990.
- [9] Spence, R. & Apperley, M. Data base navigation: An office environment for the professional. *Behaviour and Information Technology* 1 (1), 43-54, 1982.
- [10] Resnikoff, H. L. *The Illusion of Reality*. (1989) Springer-Verlag, New York.
- [11] Robertson, G. G., Card, S. K. & J. D. Mackinlay. The cognitive coprocessor architecture for interactive user interfaces. *Proceedings of the ACM SIGGRAPH Symposium on User Interface Software and Technology*, (Williamsburg, Virginia), November 1989, 10-18.
- [12] Robertson, G. G., Mackinlay, J. D., & Card, S. K. Cone Trees: Animated 3D Visualizations of Hierarchical Information. *Proceedings of CHI'91* (New Orleans, Louisiana).



MACKINLAY PLATE 1



MACKINLAY PLATE 2



MACKINLAY PLATE 3